

University of Nottingham



School of Mathematical Sciences

Topics in stochastic numerics with applications in finance and optimization

Piers Hinds

Thesis submitted to the University of Nottingham for the degree of
Doctor of Philosophy

June 2025

Abstract

This thesis is devoted to the study of two topics related to the application of numerical solution of stochastic differential equations (SDEs). First, we consider variance reduction for Lévy-driven SDEs. We develop optimality conditions for the variance reduction and propose a numerical algorithm for efficient variance reduction. The algorithm makes use of a neural network in order to approximate a control variate. Several numerical examples from option pricing are presented.

Second, we study reflected McKean-Vlasov SDEs in smooth bounded domains, for which we prove well-posedness and propagation of chaos. We also consider their application to constrained optimization problems via consensus-based optimization (CBO). We propose two CBO models, for which numerical approximation schemes correspond to constrained optimization algorithms. We test the performance of these algorithms on benchmark optimization problems as well as on an inverse problem.

Acknowledgements

I am indebted to several people whose help and encouragement have shaped this thesis. Foremost among them is my supervisor, Professor Mikhail Tretyakov, whose immense kindness and patience I am very grateful for. It has been a pleasure to work with you on a range of topics over the last five years, and I will miss it greatly when I leave.

I am also grateful to Dr. Akash Sharma for his help and guidance, particularly at the start of my PhD. Beginnings are delicate times, and having someone to answer my questions has made a huge difference to my PhD experience. I would also like to thank Professor Alain-Sol Sznitman and Professor Terry Lyons for useful discussions and suggestions.

Finally, I owe my thanks to the friends I have made at Nottingham - Alice, Emily, Erroxé, Ewan, Josh, Liam, Michael, Niamh, and Sonia - for their friendship and support throughout. I am especially appreciative of Liam and Alice, without whom my experience would have been significantly less joyful.

Contents

1	Introduction	1
1.1	Notation	3
2	Preliminaries	5
2.1	Probability and stochastic processes	5
2.2	Stochastic numerics	10
3	Variance reduction for stochastic differential equations	13
3.1	Introduction	13
3.2	Preliminaries	18
3.2.1	Variance reduction	18
3.2.2	Lévy processes	22
3.2.3	Artificial Neural Networks	29
3.3	Variance reduction for stochastic differential equations . .	30
3.3.1	Control Variates	32
3.3.2	Importance sampling	33
3.3.3	Combining method	34
3.3.4	Numerical algorithm	36
3.4	Lévy-driven stochastic differential equations	40
3.4.1	Numerical algorithm	47
3.5	Numerical experiments	50
3.5.1	Discussion of deep learning setup	51
3.5.2	Diffusion models	55
3.5.3	Models with a non-singular Lévy measure	60
3.5.4	Models with a singular Lévy measure	62
3.5.5	Comparison with alternative complexity and variance reduction methods	65
3.5.6	Transfer learning	66
3.6	Conclusion and future research	67
4	Well-posedness and approximation of reflected McKean-Vlasov	

SDEs	69
4.1 Introduction	69
4.2 Preliminaries	72
4.2.1 Reflected stochastic differential equations	72
4.2.2 McKean-Vlasov stochastic differential equations	75
4.2.3 Consensus-based optimization	76
4.3 Reflected McKean-Vlasov stochastic differential equations	79
4.3.1 Reflected consensus-based models for optimization	81
4.4 Well-posedness results	84
4.4.1 Well-posedness of reflected McKean-Vlasov SDEs	86
4.4.2 Well-posedness of the particle system	95
4.4.3 Well-posedness of the CBO models	96
4.5 Convergence of interacting particle system to the mean-field limit	97
4.6 Convergence of mean-field limit of CBO-type models to the global minimum	107
4.7 Numerical tests	116
4.7.1 Approximation of the interacting particle system	118
4.7.2 Ackley function	119
4.7.3 Non-convex function with heart-shaped constraint	120
4.7.4 High-dimensional Rastrigin function	122
4.7.5 Testing repelling CBO on Rosenbrock function	123
4.7.6 Inverse problem	125
4.7.7 Non-Lipschitz function	127
4.8 Conclusion and future research	128
5 Conclusion	130

1 Introduction

Stochastic differential equations (SDEs) are used to describe the time-evolution of random systems in several fields including mathematical finance, physics, biology, optimization and sampling, and machine learning [136, 72, 2, 28, 127]. The numerical solution of such equations has made them a practical modelling tool, in the sense that we can numerically simulate these models with a desired level of accuracy [101, 80]. In this thesis, two distinct topics related to the numerical solution of SDEs are studied; the first with an application to financial option pricing and the second with application to constrained optimization.

The first topic, which is contained in Chapter 3, concerns variance reduction methods for Lévy-driven SDEs. Such methods naturally arise when using the Monte Carlo (MC) technique to compute the expectation of functionals of solutions to SDEs [101]. In finance, the computation of these expectations corresponds to finding the current prices of financial derivatives and their sensitivities [136]. Methods of variance reduction play an important role in minimizing the error introduced by the MC method, thereby improving the computational efficiency of the MC simulation. Additionally, Lévy-driven SDEs correspond to financial models in which the market observable's trajectory exhibits discontinuous behaviour. Such behaviour is observed in markets, and considerable study has therefore been devoted to these models (see [133] and references therein).

The contribution of our work is the introduction of a variance reduction algorithm, which uses deep learning to learn an 'optimal' control variate that, in theory, yields a zero-variance MC estimator. We develop optimality conditions for the variance reduction, which demonstrates that it is possible to achieve perfect variance reduction. Our work builds on [139, 98], but extends these works to the Lévy-driven case. We discuss the contributions further in Chapter 3. We present several numerical examples related to financial derivative pricing, which demonstrate the

practical effectiveness of the algorithms.

The second topic, which is contained in Chapter 4, concerns McKean-Vlasov SDEs confined to a bounded domain via a reflection mechanism. McKean-Vlasov SDEs arise as the limit of interacting particle systems and have found application in statistical physics, mean-field games, and systematic risk modelling [70, 93, 86, 27]. At the same time, reflected SDEs are used to model processes constrained to a domain, where the reflection mechanism, which forces the process to stay within the domain, is encoded within the dynamics. We study the combination of these models, which has direct application to constrained consensus-based optimization (CBO). CBO, first considered in [111], is a metaheuristic global optimization algorithm in which interacting particles explore and exploit the domain of an objective function. CBO has attracted considerable interest in recent years due to having a strong analytical foundation, being derivative-free and being practically effective on a range of benchmark problems [28, 30, 43, 134, 71]. More recently, there has been interest in deploying CBO models in optimization problems with constraints [6, 18, 31]; however, it was not until our work [65] (see also [8], which was developed independently at the same time) that the reflection mechanism for the constraints was encoded in the continuous-time SDE via reflected SDEs. This preserves the rigorous analytical foundation of CBO models and allows us to prove convergence results.

The contribution of our work is as follows. We prove well-posedness of the reflected McKean-Vlasov SDE and its corresponding particle system in smooth bounded domains. Additionally, we prove convergence of the particle system to the mean-field limit as the number of particles goes to infinity, known as the propagation of chaos. Furthermore, we introduce two constrained CBO models as specific examples of reflected McKean-Vlasov SDEs, and prove that these converge to the minimizer of a global optimization problem. We demonstrate the effectiveness of these models on several benchmark constrained optimization problems, including an inverse problem.

The work in this thesis has resulted in the following publications:

- Hinds, P.D., and Tretyakov, M.V. "Neural variance reduction for stochastic differential equations". *J. Comp. Finance* 27(3) (2023), pp. 1–41.
- Hinds, P. D., Sharma, A., and Tretyakov, M. V. "Well-posedness and approximation of reflected McKean-Vlasov SDEs with applications". *Mathematical Models and Methods in Applied Sciences* 35(08) (2025), pp. 1845–1887.

Additionally, parts of this work I presented at the Research Students' Conference in Probability and Statistics (RSC) in 2022 and 2023 and also at the British Applied Mathematics Colloquium (BAMC) in 2023 and 2024.

Although not presented in this thesis, research for the following two papers in preparation was also completed during my PhD:

- Garcia, J.A., Gimeno, R., Hinds, P.D., Su, H., and Tretyakov, M.V. "Understanding inflation: insights from the term structure of inflation risks". (2025).
- Garcia, J.A., Gimeno, R., Hinds, P.D., Su, H., and Tretyakov, M.V. "The return of inflation and inflation risks". (2025).

These works are on making use of the market prices of inflation options to evaluate risk-neutral densities for inflation expectations.

This thesis is structured as follows. In Chapter 2 we introduce preliminary material on probability, stochastic processes, and stochastic numerics. The variance reduction for SDEs project is presented in Chapter 3 while Chapter 4 is dedicated to reflected McKean-Vlasov SDEs. Some concluding remarks are written in Chapter 5.

1.1 Notation

For two vectors $x, y \in \mathbb{R}^d$ we use $x \cdot y$ to denote the Euclidean inner product, although sometimes we use $x^\top y$ or $\langle x, y \rangle$ when it is more conve-

nient. The associated Euclidean norm is denoted by $|x|$. Given a matrix $A \in \mathbb{R}^{d \times d}$, $\text{tr}(A)$ denotes its trace and $|A|$ denotes the Frobenius norm of A . We use $\text{Diag}(v)$ to denote the diagonal matrix with diagonal given by $v \in \mathbb{R}^d$. The identity matrix of size d is written I_d . For $r > 0$, $x \in \mathbb{R}^d$, we write $B(x, r)$ for the open ball of radius r , centred at x with respect to the Euclidean norm, $|\cdot|$. For a domain $G \subset \mathbb{R}^d$, we denote its boundary by ∂G and its closure by \bar{G} .

For a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that is twice continuously differentiable, we write $\nabla f(x)$ and $\nabla^2 f(x)$ to denote the gradient and Hessian of f at x , respectively. For a continuously differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, we denote the Jacobian of f at x by $J_f(x)$.

We use (Ω, \mathcal{F}, P) to denote a probability space, and expectation with respect to the probability measure P is denoted by $\mathbb{E}[\cdot]$. For an integrable random variable X , we denote the covariance matrix by $\text{Var}(X) := \mathbb{E}[(X - \mathbb{E}[X])(X - \mathbb{E}[X])^\top]$. The notation $\|X\|_p$ denotes the L^p norm of X . The indicator of a measurable set $A \in \mathcal{F}$ is written $\mathbb{1}_A$ or I_A .

Throughout the text, the symbol $C > 0$ denotes a finite positive constant whose value may change from line to line but never depends on the main parameters under consideration.

2 Preliminaries

This chapter serves as an introduction to some background material needed for both Chapter 3 and Chapter 4. Both of these chapters provide their own additional preliminaries. Here, we introduce background material on probability, stochastic processes, and the numerical approximation of solutions to stochastic differential equations. For further reading, we refer the reader to [123, 15] for probability, [118] for stochastic processes, [108] for stochastic differential equations (SDEs), and [101] for stochastic numerics.

2.1 Probability and stochastic processes

Definition 2.1. A **probability space** is a triple (Ω, \mathcal{F}, P) consisting of:

- a set Ω , called the *sample space*;
- a σ -algebra \mathcal{F} on Ω , called the *family of events*;
- a measure P on \mathcal{F} with $P(\Omega) = 1$, called the *probability or probability measure*.

A probability space is called **complete** if for any $B \in \mathcal{F}$ with $P(B) = 0$, all subsets $A \subset B$ are included in \mathcal{F} (and consequently assigned measure zero). Given a probability space, one can always extend it to a complete probability space. As such, we do not lose any generality working only with complete spaces, and it is a canonical assumption (usually part of the so-called *usual conditions*) in the theory of continuous-time stochastic processes [118]. Henceforth, we work exclusively with complete probability spaces.

Definition 2.2. A (\mathbb{R}^d -valued) **random variable** on a probability space (Ω, \mathcal{F}, P) is a measurable function $X : \Omega \rightarrow \mathbb{R}^d$ from the sample space to state space \mathbb{R}^d , where \mathbb{R}^d is equipped with the Borel σ -algebra.

The above definition can be extended to define an E -valued random variable given a measurable space (E, \mathcal{E}) as the state space; however, we

only make use of \mathbb{R}^d -valued random variables. When $d = 1$, we will call them **real-valued random variables**.

For a random variable, X , we define the L^p -norm of X by

$$\|X\|_p = (\mathbb{E}|X|^p)^{1/p}, \quad (2.1)$$

for $0 < p < \infty$, and we let $L^p(\Omega, \mathcal{F}, P)$ denote the set of random variables with

$$\|X\|_p < \infty. \quad (2.2)$$

Definition 2.3. A **stochastic process** is a collection of random variables, $X = (X(t), t \in \mathbb{T})$, indexed by an index set \mathbb{T} .

The index set of a stochastic process typically represents time. If $\mathbb{T} = \mathbb{N}$, we call the stochastic process *discrete-time* and if $\mathbb{T} = [0, \infty)$ (or some interval subset of $[0, \infty)$), we call it *continuous-time*. We will assume that we are working with a continuous-time stochastic process unless we specify that it is discrete-time.

A stochastic process can be thought of as a mapping from the joint space $(\Omega \times [0, \infty), \mathcal{F} \otimes \mathcal{B}([0, \infty)))$ to $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. We say that the process is **measurable** if this map is measurable. We will alternate between the conventional notation $X(t, \omega)$ and $X(t)$, depending on what we would like to emphasize. We may abuse notation and write $X(t)$ to describe the whole process when we mean $(X(t), t \in \mathbb{T})$. For a fixed sample point $\omega \in \Omega$, the corresponding **sample path** of a stochastic process, X , is the map $t \mapsto X(t, \omega)$.

Additional structure can be added to a probability space to describe information flow over time.

Definition 2.4. A **filtration** on a probability space (Ω, \mathcal{F}, P) is a family of σ -algebras, $(\mathcal{F}_t)_{t \geq 0}$, such that:

- for each $t \geq 0$, \mathcal{F}_t is a sub- σ -algebra of \mathcal{F} ;
- for $s \leq t$, $\mathcal{F}_s \subseteq \mathcal{F}_t$.

A filtration is said to satisfy the **usual conditions** if it is complete and right-continuous. Any filtration can be refined into a complete filtration and, as the name suggests, this is a standard assumption in the study of continuous-time stochastic processes [118]. We assume any filtration we work with satisfies the usual conditions.

A probability space (Ω, \mathcal{F}, P) equipped with a filtration $\mathbb{F} := (\mathcal{F}_t)_{t \geq 0}$ is called a **filtered probability space** and is denoted by the quadruple $(\Omega, \mathcal{F}, \mathbb{F}, P)$. At times, we may abuse notation and write $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$ instead.

Let $(X(t), t \geq 0)$ be a stochastic process on $(\Omega, \mathcal{F}, \mathbb{F}, P)$. We say that the process is **adapted** to \mathbb{F} , if for every $t \geq 0$, the random variable $X(t)$ is \mathcal{F}_t -measurable.

A stochastic process, $(X(t), t \geq 0)$, induces a filtration via the prescription $\mathcal{F}_t = \sigma(X(s) : 0 \leq s \leq t)$ for each $t \geq 0$, where $\sigma(Y)$ denotes the σ -algebra generated by a random variable Y (the smallest σ -algebra such that Y is measurable). This filtration is called the **natural filtration** of the process $(X(t), t \geq 0)$. Any process is adapted to its natural filtration.

Definition 2.5. A **Wiener process** (or **Brownian motion**) on a probability space (Ω, \mathcal{F}, P) is a continuous-time stochastic process, $(W(t), t \geq 0)$, with the following properties:

1. $W(0) = 0$;
2. W has independent increments;
3. For any $0 \leq s < t$, the increment $W(t) - W(s) \sim \mathcal{N}(0, t - s)$;
4. The sample paths of W are continuous.

Let us note that we will use the names Wiener process and Brownian motion interchangeably throughout. We can define a d -dimensional Wiener process, $(W(t), t \geq 0)$ as

$$W(t) = (W_1(t), \dots, W_d(t))^\top, \quad (2.3)$$

where $W_i, i = 1, \dots, d$, are independent one-dimensional Wiener processes.

For a Brownian motion $(W(t), t \geq 0)$ (for now one-dimensional) and a progressively measurable stochastic process $(X(t), t \geq 0)$ which is adapted to the natural filtration of W and satisfies

$$\mathbb{E} \int_0^T X^2(t) dt < \infty, \quad (2.4)$$

one can construct the Ito integral:

$$\int_0^t X(s) dW(s), \quad (2.5)$$

for $t \in [0, T]$ (see e.g. [108]).

The Ito integral in (2.5) can itself be considered as a stochastic process in t and, analogously to (Riemann) integral equations, one can consider the case when the integrand depends on the process itself. For example, given a suitably regular function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, we can consider the (Ito) integral equation:

$$X(t) = X(0) + \int_0^t \sigma(X(s)) dW(s), \quad (2.6)$$

where $X(0)$ is specified. Equations of the form (2.6) are called **stochastic differential equations (SDEs)** and are often written in differential form:

$$dX(t) = \sigma(X(t)) dW(t), \quad (2.7)$$

which is purely an abbreviated version of (2.6), provided $X(0)$ is still specified. More generally, we can extend to the multidimensional case with time-dependent coefficients. Given a d -dimensional Brownian motion $(W(t), t \geq 0)$ and $b : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, we can write the SDE

$$X(t) = X(t_0) + \int_{t_0}^t b(s, X(s)) ds + \int_{t_0}^t \sigma(s, X(s)) dW(s), \quad (2.8)$$

$t \geq t_0$, with $X(t_0)$ given. The first integral in (2.8) is an ordinary Riemann integral and the multidimensional Ito integral should be interpreted as

$$\left(\int_{t_0}^t \sigma(s, X(s)) dW(s) \right)_i = \sum_{j=1}^d \int_{t_0}^t \sigma_{ij}(s, X(s)) dW_j(s). \quad (2.9)$$

The (time-dependent) vector field b is often referred to as the *drift* of the process X and the matrix-valued function σ is often called the *volatility* or *diffusion coefficient* of X .

Existence and uniqueness of a (strong) solution to the SDE (2.8) depends on the functions b and σ . A sufficient condition for existence and uniqueness is that b and σ are globally Lipschitz with at most linear growth at infinity (see e.g. [108]). Weaker conditions can be considered, see for example [73, 76].

Assume that b and σ satisfy these conditions and let X be the unique strong solution to (2.8). The process X is a Feller process and has the generator [53]

$$L_t u(x) := \frac{1}{2} \text{tr}[\sigma(t, x) \sigma^\top(t, x) \nabla^2 u(x)] + \langle b(t, x), \nabla u(x) \rangle. \quad (2.10)$$

Fix $0 \leq t_0 < T < \infty$, and for any $t_0 \leq t \leq T$ and any $x \in \mathbb{R}^d$, let $X_{t,x}$ be the solution of (2.8) satisfying the initial condition $X_{t,x}(t) = x \in \mathbb{R}^d$. For a regular enough test function, f , consider the following expectation as a function of the initial condition:

$$u(t, x) = \mathbb{E}[f(X_{t,x}(T))]. \quad (2.11)$$

The **Feynman-Kac formula** [79, 73] states that the function u satisfies the following Cauchy problem for the parabolic PDE:

$$\frac{\partial u}{\partial t} + L_t u = 0, \quad (t, x) \in [t_0, T) \times \mathbb{R}^d \quad (2.12)$$

$$u(T, x) = f(x), \quad x \in \mathbb{R}^d. \quad (2.13)$$

If, instead of (2.12), we would like to consider the general linear

parabolic PDE

$$\frac{\partial u}{\partial t} + L_t u + c(t, x)u + g(t, x) = 0, \quad (2.14)$$

for $t \in [t_0, T)$, $x \in \mathbb{R}^d$ and with the same terminal condition (2.13), we would consider the function

$$u(t, x) = \mathbb{E} \left[f(X_{t,x}(T)) e^{\int_t^T c(s, X_{t,x}) ds} + \int_t^T g(s', X_{t,x}) e^{\int_t^{s'} c(s, X_{t,x}) ds} ds' \right] \quad (2.15)$$

instead. The Feynman-Kac formula (2.15) allows us to solve problems which are deterministic in nature, i.e. the Cauchy problem (2.14) and (2.13), with probabilistic methods via the computation of the expectation (2.15). Equivalently, if one wishes to compute the expectation (2.15), the Feynman-Kac formula allows us to solve the related PDE.

2.2 Stochastic numerics

Given a stochastic process $(X(t), t \geq 0)$ defined as the strong solution to

$$X(t) = X(0) + \int_0^t b(s, X(s)) ds + \int_0^t \sigma(s, X(s)) dW(s), \quad (2.16)$$

with $X(0)$ given, one would like to sample the process X , either at a certain time or entire trajectories. For example, to compute the expectation in (2.15) via the Monte Carlo method, one would need to sample the random variable $X(T)$. In general, an analytical solution to (2.16) cannot be found, but the solution can be approximated via a (stochastic) numerical integration scheme.

A typical way to construct an approximation to the solution of (2.16) is to simulate sample paths of a discrete-time stochastic process such that at each time-point in the discretization, the approximate random variable is close, in some sense, to the true solution. Consider the interval $[0, T]$ and let $h = T/N$ for some $N \in \mathbb{N}$. Consider the uniform discretization of

this interval:

$$t_0 = 0, \quad (2.17)$$

$$t_k = t_0 + kh, \quad k = 1, \dots, N. \quad (2.18)$$

We approximate the solution X with the discrete-time process $(X_n^h, n = 1, \dots, N)$, which is constructed by a specified numerical integration scheme. We say that a numerical integration scheme converges in mean-square with order q if there exists a constant $C > 0$, which does not depend on h or k , such that

$$(\mathbb{E}|X(t_k) - X_k^h|^2)^{\frac{1}{2}} \leq Ch^q. \quad (2.19)$$

In many applications, however, we only care about closeness in distribution. A numerical integration scheme converges in the weak-sense at time T with order q if there exists $C > 0$, which does not depend on h , such that

$$|\mathbb{E}f(X(T)) - \mathbb{E}f(X_N^h)| \leq Ch^q, \quad (2.20)$$

for f from a sufficiently large family of functions [101].

The **Euler scheme** (sometimes Euler-Maruyama) [92] is the simplest numerical integration scheme. The Euler approximation is constructed as follows:

$$X_0^h = X_0, \quad (2.21)$$

$$X_{k+1}^h = X_k^h + b(t_k, X_k^h)h + \sigma(t_k, X_k^h)Z_{k+1}, \quad (2.22)$$

where $Z_k \sim \mathcal{N}(0, h)$, $k = 1, \dots, N$ are independent and identically distributed. The Euler scheme has mean-square order $1/2$ [54, 101] and weak order 1 [97, 101]. If the random variables $\{Z_k\}$ in (2.22) are replaced with

$$Z_k = \begin{cases} +\sqrt{h} & p = \frac{1}{2}, \\ -\sqrt{h} & p = \frac{1}{2}, \end{cases} \quad (2.23)$$

$k = 1, \dots, N$, which are independent, we obtain the crude Euler method,

which also has weak order 1 (e.g. [101, 95]). For a higher rate of mean-square convergence, Milstein's method can be considered [101]. In one-dimension, this is defined by the update rule:

$$X_{k+1}^h = X_k^h + b(t_k, X_k^h)h + \sigma(t_k, X_k^h)Z_{k+1} \quad (2.24)$$

$$+ \frac{1}{2}\sigma(t_k, X_k^h)\frac{\partial\sigma}{\partial x}(t_k, X_k^h)[Z_{k+1}^2 - h], \quad (2.25)$$

where $Z_k \sim \mathcal{N}(0, h)$, $k = 1, \dots, N$ are independent and identically distributed. Milstein's method has mean-square order 1 and weak order 1 under certain regularity conditions [101].

3 Variance reduction for stochastic differential equations

In this chapter, we study variance reduction for Monte Carlo (MC) integration of functionals of the solutions to stochastic differential equations (SDEs). We propose an algorithm, in which we approximate (theoretically optimal) control variates with neural networks, to obtain a reduced-variance estimator for the MC integration. We consider SDEs driven by Lévy processes, for which we extend some theoretical results and present several numerical examples. This chapter is based on the paper [66]:

- Hinds, P.D., and Tretyakov, M.V. "Neural variance reduction for stochastic differential equations". *J. Comp. Finance* 27(3) (2023), pp. 1–41.

but has significantly more exposition and some additional material.

This chapter is structured as follows. Section 3.1 serves as an introduction to the topic and a review of related literature is presented. Section 3.2 contains preliminary material related to variance reduction methods and also to Lévy-driven SDEs. This builds on the preliminary material of Chapter 2 and is specific to this chapter. In Section 3.3, we consider the case where SDEs are driven purely by Brownian motion and the corresponding PDE problem. We recall the optimality conditions for variance reduction and present a numerical algorithm to find an approximation of the optimal control variate using deep learning. In Section 3.4, we consider the more general case where SDEs are driven by Lévy noise. We derive optimality conditions for variance reduction and present a corresponding numerical algorithm. Finally, Section 3.5 contains several numerical examples from mathematical finance.

3.1 Introduction

SDEs driven by Lévy processes arise as a modelling tool in several fields including finance, insurance, physics, and chemistry [34, 84, 72]. In

finance, SDEs are employed to model the stochastic dynamics of market observables (asset prices, interest rates, FX rates, etc.), with the goal of pricing and hedging financial derivatives written on these market observables (see e.g. [23, 34, 57, 136] and references therein). In particular, SDEs are used to find current prices and hedges of options which are not traded (or not regularly traded on exchanges) and to evaluate future option prices for risk management purposes. The task of pricing an option is, fundamentally, the computation of an expectation. Indeed, in many cases including option pricing, one wants to compute an expectation of the form

$$u(t, x) = \mathbb{E} \left[f(X_{t,x}(T)) \right], \quad (3.1)$$

where $X_{t,x}(s)$, $s \geq t$, is the solution to a system of SDEs with the initial condition $X_{t,x}(t) = x$, $T > 0$ is a terminal (maturity) time, and f is a real-valued function (payoff in the case of option pricing).

The most general way of approximating the expectation (3.1) is via weak-sense numerical integration of the SDE (in order to produce samples of the random variable), coupled with Monte Carlo (MC) integration to approximate the expectation [57, 101, 112]. When applying this probabilistic method, two sources of error are present in the numerical approximations: the error arising from numerical integration of the SDE and the error from approximating the expectation with a MC estimate. Methods of variance reduction, which we study in the sequel, play an important role in minimizing the MC error. The focus of this chapter is on the implementation of practical variance reduction methods in this setting.

As it is well known (see e.g. [42, 46]), the function $u(t, x)$ solves the related, via the Feynman-Kac formula, partial differential equation (PDE) problem in the diffusion case or partial integro-differential equation (PIDE) problem in the case of general Lévy-driven SDEs. Consequently, when one wishes to solve a P(I)DE problem, the solution $u(t, x)$ can be approximated numerically via the probabilistic method described previously. In cases where the spatial dimension (say d) is large, the

MC method becomes particularly effective, since traditional grid-based methods become computationally infeasible, due to the complexity of the discretization scaling exponentially with d . In contrast, the complexity of the MC method is independent of d and the complexity of the numerical integration of the SDEs scales linearly with d . The price to pay for this better scaling, however, is that when using the probabilistic method, we solve the P(I)DE problem at only one point, say (s, x_0) , rather than recovering the global solution.

When SDEs are driven purely by Brownian motion, i.e. when X is the solution to

$$dX = b(t, X)dt + \sigma(t, X)dW, \quad (3.2)$$

there is an optimality condition (see [105, 98, 101] or Theorem 3.2 here) relating the optimal variance reduction with the solution of the PDE problem, $u(t, x)$, and its spatial derivatives, $\partial u(t, x)/\partial x^i$. Of course, if the PDE solution $u(t, x)$ is known, there would be no need to resort to probabilistic methods in the first place. The importance of this theoretical result lies in the fact that it demonstrates that perfect variance reduction is possible and, if an approximation of the solution $u(t, x)$ can be learned, it can induce efficient variance reduction. In addition, this variance reduction method does not bias the MC approximation (see e.g. [105, 100, 101]), indicating that the approximation of the required $u(t, x)$ and its derivatives need not have a high accuracy in order to be useful.

To make variance reduction practical, a suitable method of constructing $u(t, x)$ and its spatial derivatives should be inexpensive. Hence, a trade-off between accuracy and computational costs in finding $u(t, x)$ and $\partial u(t, x)/\partial x^i$ is needed. To address this problem, it was suggested in [100] (see also [101] and [12]) to exploit conditional probabilistic representations of $u(t, x)$ in conjunction with linear regression, which allows one to evaluate $u(t, x)$ and $\partial u(t, x)/\partial x^i$ using the single auxiliary set of approximate trajectories starting from an initial position. This leads to obtaining sufficiently inexpensive, but useful for variance reduction, estimates of $u(t, x)$ and $\partial u(t, x)/\partial x^i$. The drawback of this approach is the reliance on

linear regression. To make it computationally feasible, a careful selection of basis functions for linear regression is needed. This selection is heavily problem dependent and limits the applicability of this variance reduction approach.

In this direction, Vidales, Šiška and Szpruch [139] propose several algorithms with the aim of constructing a control variate via deep learning. An analogue of the linear regression algorithm from [100] is proposed, using a deep neural network (rather than linear regression) to approximate the PDE solution. The constructed rough solution will not necessarily be accurate, but no bias is introduced due to the use of the MC technique. In addition to this, a method is put forward to learn the control variates directly, without first constructing a solution to the PDE [139, Algorithm 4]. In order to do this, the empirical variance is used as an objective (loss) function and minimized, in the standard way, using a stochastic gradient algorithm. In all cases, the training of neural networks is carried out offline over a parametric family of SDEs (PDEs). In the context of financial option pricing, this means that training one network is sufficient for a single model-payoff combination.

The key objective of this chapter is to show that deep learning (i.e. nonlinear regression) can successfully replace linear regression in variance reduction for SDEs in a universal manner, i.e. without need to do tuning for each new system of SDEs which is important for applications, in particular for those arising in financial engineering. The use of neural networks instead of linear regression as in [100] eliminates the need of finding a specific set of basis functions, thus making the variance reduction truly practical.

Furthermore, in contrast to [139] and many other deep learning works related to computational finance and more generally to efficient SDEs and PDEs simulations (see e.g. [124, 63, 52, 129] and references therein), here we do not rely on an (often costly) offline training of neural networks but rather we do network training online together with actual simulation of the required expectation $u(t, x)$ with reduced variance. In

other words, we propose a black-box fashion practical variance reduction tool which does not require any lengthy pre-training and tuning, analogously to how the variance reduction method of [100] intended to work but without the limitations due to linear regression. Note that this approach differs from [139], where training is done offline across a parametric family of PDEs.

For general Lévy-driven SDEs, variance reduction has been considerably less studied than in the Brownian case. To the best of our knowledge, a result demonstrating that there exists an optimal choice of control variates, analogous to the results (in the diffusion case) in [105, 98] (see also [101]), was first obtained in our work [66]. We note that in the context of financial applications, there have been several effective approaches towards variance reduction under Lévy-driven models (e.g. [122, 39, 121] and references therein).

We also remark that there are other techniques than variance reduction aimed at reducing computational complexity of MC simulations: multi-level MC method [55] and quasi-Monte Carlo method [38]. They have their own areas of applicability and can potentially be combined with deep learning and neural SDEs considered in this paper. Here, we restrict ourselves to considering the use of deep learning for improving the plain-vanilla MC technique via variance reduction.

The work in this chapter is most closely related to [98, 100, 139]. We show that for Lévy-driven SDEs, there exist conditions ensuring the optimal variance reduction, which is an analogous result to the one in the Brownian SDEs case [98]. We also provide details of a novel algorithm which produces low variance simulations of a given system of SDEs. Given a system of SDEs, we introduce control variates and parameterize them with neural networks. This can be interpreted in two ways. First, we can treat the problem of fitting the neural network as a nonlinear regression problem, for which we need to solve an optimization problem. As such, the proposed method can be seen as a more general framework than that proposed in [100]. Second, the use of a neural

network as a coefficient in SDEs gives rise to neural SDEs (see e.g. [138, 67]), which is essentially a generative model (see [77]). Regardless of the perspective, the parameters of the neural SDEs can be learned from numerical simulations of the system, in such a way that the empirical variance is minimized. SDEs' trajectories with high accuracy can then be simulated using the approximate control variates. We leverage the fact that no prescribed accuracy is required for the approximated control variates, since they do not introduce bias into the MC approximation. This algorithm is not limited to the case that SDEs are driven only by Brownian motion, as in [100, 139], and can be applied in the general setting of Lévy noise, even when we have infinite activity of jumps.

3.2 Preliminaries

In this section, we provide some background material specific to this chapter. We recall the Monte Carlo method and motivate the need for variance reduction techniques. We describe two well-known methods of variance reduction and explain how these methods can be applied in the context of SDEs. We also include some preliminaries for Lévy processes and, in particular, Lévy-driven SDEs. Finally, we include a brief description of artificial neural networks.

3.2.1 Variance reduction

Let (Ω, \mathcal{F}, P) be a probability space and let $\xi \in L^2(\Omega, \mathcal{F}, P)$ be an arbitrary real-valued random variable for which we would like to compute

$$\mu := \mathbb{E}[\xi]. \quad (3.3)$$

This expectation can be approximated by the MC estimator

$$\hat{\mu}_M := M^{-1} \sum_{m=1}^M \xi_m, \quad (3.4)$$

where $\{\xi_m : m = 1, \dots, M\}$ are independent copies of ξ [24]. The estimator $\hat{\mu}_M$ is unbiased (for any M , $\mathbb{E}\hat{\mu}_M = \mu$), and the strong law of large numbers ensures that it converges almost surely to μ . In practical applications, it is important to understand how close the MC estimator is to μ for a finite M . The MC error, ϵ_M , is defined to be the difference between the MC estimator and μ

$$\epsilon_M := \hat{\mu}_M - \mu. \quad (3.5)$$

The Central Limit Theorem implies that

$$\frac{\sqrt{M}\epsilon_M}{\sqrt{\text{Var}[\xi]}} \stackrel{M \rightarrow \infty}{\approx} \mathcal{N}(0, 1), \quad (3.6)$$

and so, for large enough M , we have the approximation

$$\epsilon_M \sim \mathcal{N}\left(0, \frac{\text{Var}[\xi]}{M}\right). \quad (3.7)$$

As a result, we can report confidence intervals of the MC error in addition to the MC estimate. Asymptotic $(1 - \delta)$ -confidence intervals are given as

$$\hat{\mu}_M \pm z_{1-\delta/2} \sqrt{\frac{\text{Var}[\xi]}{M}}, \quad (3.8)$$

where z_δ denotes the δ quantile of the standard normal distribution. The root mean-squared error is of order $M^{-1/2}$ which, although slow, is dimension independent. The root mean-squared error also depends on $\sqrt{\text{Var}[\xi]}$, the standard deviation of the random variable whose mean we are estimating.

The objective of variance reduction techniques is to lessen this coefficient before sampling. Of course, this does nothing to change the order of convergence - it remains $\mathcal{O}(M^{-1/2})$ - however, with a smaller coefficient we require fewer samples (less computational effort) in order to achieve a prescribed tolerance level for the MC error.

Variance reduction techniques, therefore, can be thought of as finding

another random variable, say $\tilde{\xi}$, with

$$\begin{aligned}\mathbb{E}[\tilde{\xi}] &= \mathbb{E}[\xi], \\ \text{Var}[\tilde{\xi}] &\ll \text{Var}[\xi],\end{aligned}$$

and then sampling from $\tilde{\xi}$ instead of ξ . The construction of the random variable $\tilde{\xi}$ is determined by the particular variance reduction technique. In general, this construction can be highly problem dependent [57]. We briefly describe the two general and well-known variance reduction techniques of *control variates* and *importance sampling*.

Control variates

The technique of control variates relies on an additional auxiliary random variable which has zero mean and is correlated with ξ [57]. Formally, given $\zeta \in L^2(\Omega, \mathcal{F}, P)$ with $\text{Var}[\zeta] > 0$ and

$$\begin{aligned}\mathbb{E}[\zeta] &= 0, \\ \mathbb{E}[\xi\zeta] &\neq 0,\end{aligned}$$

we construct the random variable $\tilde{\xi}$ as

$$\tilde{\xi} := \xi - \frac{\mathbb{E}[\xi\zeta]}{\text{Var}[\zeta]}\zeta. \quad (3.9)$$

The variance of this new random variable is then

$$\text{Var}[\tilde{\xi}] = \text{Var}[\xi] - \frac{(\mathbb{E}[\xi\zeta])^2}{\text{Var}[\zeta]}, \quad (3.10)$$

which is necessarily less than $\text{Var}[\xi]$. In fact, the coefficient $-\frac{\mathbb{E}[\xi\zeta]}{\text{Var}[\zeta]}$ in (3.9) is optimal in the sense that the variance of the random variable

$$\xi_\lambda := \xi + \lambda\zeta, \quad (3.11)$$

$\lambda \in \mathbb{R}$, is minimized with the choice $\lambda = -\frac{\mathbb{E}[\xi\zeta]}{\text{Var}[\zeta]}$. To see this, note that

$$\text{Var}[\xi_\lambda] = \lambda^2 \text{Var}[\zeta] + 2\lambda \mathbb{E}[\xi\zeta] + \text{Var}[\xi], \quad (3.12)$$

which is a quadratic in λ and hence minimized at $\lambda = -\frac{\mathbb{E}[\xi\zeta]}{\text{Var}[\zeta]}$.

Suppose we have a random variable $\tilde{\zeta}$, whose expectation is not zero, but is correlated with ξ . Then we can construct a control variate by centring the random variable

$$\zeta = \tilde{\zeta} - \mathbb{E}\tilde{\zeta}. \quad (3.13)$$

Of course, to utilize this control variate, we need to know the expectation $\mathbb{E}\tilde{\zeta}$. For example, we may know the expectation analytically, or we may have a numerical method to acquire it (in which case this numerical method should be substantially quicker than MC).

Importance sampling

The technique of importance sampling involves constructing a new random variable by multiplying ξ by a Radon-Nikodym derivative and then sampling from the associated measure in order to keep the expectation fixed [57]. Formally, let Q be a probability measure on (Ω, \mathcal{F}) , such that P is absolutely continuous with respect to Q (that is, if $Q(A) = 0$ then $P(A) = 0$). Then we can write

$$\mathbb{E}_P[\xi] = \int_{\Omega} \xi(\omega) dP(\omega) \quad (3.14)$$

$$= \int_{\Omega} \xi(\omega) \frac{dP}{dQ}(\omega) dQ(\omega), \quad (3.15)$$

where $\frac{dP}{dQ} \geq 0$ denotes the Radon-Nikodym derivative of P with respect to Q . Writing $\tilde{\xi} = \xi \frac{dP}{dQ}$, we have

$$\mathbb{E}_P[\xi] = \mathbb{E}_Q[\tilde{\xi}], \quad (3.16)$$

establishing that we can draw samples of $\tilde{\xi}$ under the measure Q to construct our MC estimator. The measure Q is arbitrary and, to be useful, should be chosen such that $\tilde{\xi}$ has smaller variance (under Q) than ξ (under P).

3.2.2 Lévy processes

In this section, we present some background theory on Lévy processes. The purpose of this section is to establish the notion of SDEs driven by Lévy processes, which are a useful tool for financial modelling [133] among other applications [84, 72]. Lévy SDEs will be the central object in the Section 3.4. For the material in this section, we follow [3].

We begin with the definition of a Lévy process.

Definition 3.1. A **Lévy process** is a continuous-time stochastic process, $(X(t), t \geq 0)$ with the following properties:

- L1. $X(0) = 0$ almost surely;
- L2. (Independent increments) For any $n \in \mathbb{N}$ and any $0 \leq t_1 < \dots < t_n$, the increments $X(t_2) - X(t_1), \dots, X(t_n) - X(t_{n-1})$ are independent;
- L3. (Stationary increments) For any $t \geq 0, h > 0$, the increment $X(t+h) - X(t)$ is equal in distribution to $X(h)$;
- L4. (Stochastic continuity) For any $t \geq 0, X(s) \xrightarrow{P} X(t)$ as $s \rightarrow t$ (that is, for any $\varepsilon > 0$, we have $\lim_{h \rightarrow 0} P(\|X(t+h) - X(t)\| > \varepsilon) = 0$).

The property (L4) ensures that there exists a (right-continuous with left limits) modification of the process. Therefore, we will assume that any Lévy process we work with is indeed càdlàg. Stochastic continuity is a much weaker assumption than assuming almost surely continuous sample paths, and, as a result, Lévy processes can exhibit jumps (discontinuities of the sample paths). Due to (L4), the probability of a jump at any $t \geq 0$ is 0; however, over any interval $[0, t]$ these null events can add up to something substantial. Working with a càdlàg modification is desirable (some authors consider this as part of the definition, see e.g. [84]) for

several reasons; one of which is that the number of discontinuities of a càdlàg path is at most countable.

Throughout this section, we define

$$E := \mathbb{R}^d - \{0\}. \quad (3.17)$$

The space E should be thought of as the space of possible ‘jumps’ of a Lévy process; this is the reason that 0 is omitted.

Definition 3.2. A Lévy measure, ν , is a Borel measure on E satisfying

$$\int_E (|z|^2 \wedge 1) \nu(\mathrm{d}z) < \infty \quad (3.18)$$

Lévy measures describe the intensity and distribution of jumps of Lévy processes. The condition (3.18) ensures that we only have finitely many jumps which are greater than ε in size, for any $\varepsilon > 0$. We say that a Lévy measure has **infinite activity** if $\nu(E) = \infty$ and **finite activity** if $\nu(E) < \infty$.

It is often convenient to integrate with respect to a Lévy measure over the whole space \mathbb{R}^d instead of E . This causes no mathematical change provided we add that $\nu(\{0\}) = 0$ to the definition of a Lévy measure.

Lévy processes are related, in an essential way, to infinitely divisible distributions: if $(X(t), t \geq 0)$ is a Lévy process then for any $t \geq 0$, the random variable $X(t)$ is infinitely divisible; and conversely, from any infinitely divisible distribution a Lévy process can be constructed. As a result, the Lévy-Khintchine representation [74], which gives a characterization of infinitely divisible distributions via their characteristic function, can be used to characterize a Lévy process.

If $(X(t), t \geq 0)$ is a Lévy process then the characteristic function of $X(t)$, $\phi_{X(t)}$, has the following representation:

$$\begin{aligned} \phi_{X(t)}(u) = \exp \left(t \left[i \langle b, u \rangle - \frac{1}{2} \langle u, Au \rangle \right. \right. \\ \left. \left. + \int_E [e^{i \langle u, z \rangle} - 1 - i \langle u, z \rangle \mathbb{1}_{\|z\| \leq 1}] \nu(\mathrm{d}z) \right] \right), \quad (3.19) \end{aligned}$$

where $b \in \mathbb{R}^d$, A is a symmetric, positive semidefinite $d \times d$ matrix, and ν is a Lévy measure. The triplet (b, A, ν) is uniquely determined by the distribution of the process X . Note that the integral in (3.19) is guaranteed to converge since

$$\begin{aligned} & \int_E [e^{i\langle u, z \rangle} - 1 - i\langle u, z \rangle \mathbb{1}_{\|z\| \leq 1}] \nu(\mathbf{d}z) \\ &= \int_{B(0,1) - \{0\}} [e^{i\langle u, z \rangle} - 1 - i\langle u, z \rangle] \nu(\mathbf{d}z) + \int_{\|z\| > 1} [e^{i\langle u, z \rangle} - 1] \nu(\mathbf{d}z), \end{aligned} \quad (3.20)$$

and the integrand $|e^{i\langle u, z \rangle} - 1 - i\langle u, z \rangle|$ is $\mathcal{O}(|z|^2)$ while $\nu(\{z \in E : |z| > 1\}) < \infty$ from the definition of ν .

Poisson random measures

We introduce the Poisson random measure, which is useful for describing the jump behaviour of Lévy processes.

Definition 3.3. Let μ be a σ -finite measure on $\mathcal{S} := [0, \infty) \times E$ equipped with the Borel σ -algebra. A **Poisson random measure**, N , with intensity μ is a collection of random variables $\{N(B) : B \in \mathcal{S}\}$ satisfying

1. $N(\emptyset) = 0$;
2. (σ -additivity) For any mutually disjoint $B_1, B_2, \dots \in \mathcal{S}$

$$N\left(\bigcup_{i=1}^{\infty} B_i\right) = \sum_{i=1}^{\infty} N(B_i);$$

3. (Independent scattering) For any mutually disjoint $B_1, \dots, B_n \in \mathcal{S}$, the random variables $N(B_1), \dots, N(B_n)$ are independent;
4. For $B \in \mathcal{S}$, the random variable $N(B)$ has Poisson distribution with rate $\mu(B)$.

We are interested in Poisson random measures with intensity measures of the form

$$\mu(A \times B) = \text{Leb}(A)\nu(B), \quad (3.21)$$

for $A \in \mathcal{B}([0, \infty))$, $B \in \mathcal{B}(E)$ and where Leb denotes the Lebesgue measure on $[0, \infty)$. The Poisson random measure can be thought of as a stochastic process with the prescription

$$N(t, B) := N((0, t], B), \quad (3.22)$$

for $B \in \mathcal{B}(E)$ and noting that $N(0, B) = 0$ a.s. for any $B \in \mathcal{B}(E)$.

The utility of Poisson random measures arises from the fact that any Lévy process has a corresponding Poisson random measure describing the behaviour of its jumps. For a Lévy process $X = (X(t), t \geq 0)$, define its jump process $\Delta X = (\Delta X(t), t \geq 0)$ where

$$\Delta X(t) := X(t) - X(t-).$$

Here we use the standard notation $X(t-) = \lim_{s \nearrow t} X(s)$. Then, for $A \in \mathcal{B}([0, \infty))$ and $B \in \mathcal{B}(E)$ we can define the random measure

$$N(A \times B) := |\{s \in A : \Delta X(s) \in B\}|,$$

which counts the number of jumps of a certain size in a certain time period. The random measure N is indeed a Poisson random measure with intensity $\mu(dt, dz) = dt \times \nu(dz)$, where ν is the Lévy measure of X .

Let ν be a Lévy measure and let N be a Poisson random measure with intensity $dt \times \nu(dz)$. We define the compensated Poisson random measure as

$$\hat{N}((0, t] \times B) := N((0, t] \times B) - t\nu(B). \quad (3.23)$$

for $B \in \mathcal{B}(E)$.

For a set $A \subset E$, we say that A is **bounded below** if $0 \notin \bar{A}$. If A is bounded below then

$$N(t, A) = \sum_{0 \leq s \leq t} \mathbb{1}_A(\Delta X(s)) < \infty, \quad (3.24)$$

which follows from the fact that X has càdlàg paths (see e.g. [3, Theorem

2.8.1)).

Let N be the Poisson random measure associated with a Lévy process X . Let A be bounded below and let $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be Borel measurable. We define

$$\int_A f(x)N(t, \mathbf{d}x) := \sum_{x \in A} f(x)N(t, \{x\}), \quad (3.25)$$

for every $\omega \in \Omega$. Note that the left-hand side of (3.25) is a (random) finite sum since A is bounded below. From the definition of N , $N(t, \{x\}) \neq 0$ exactly when $\Delta X_s = x$ for some $s \leq t$. Hence

$$\int_A f(x)N(t, \mathbf{d}x) = \sum_{0 \leq s \leq t} f(\Delta X_s) \mathbb{1}_A(\Delta X_s). \quad (3.26)$$

Similarly for the compensated measure \hat{N} , we define

$$\int_A f(x)\hat{N}(t, \mathbf{d}x) := \int_A f(x)N(t, \mathbf{d}x) - t \int_A f(x)\nu(\mathbf{d}x), \quad (3.27)$$

for A bounded below. The compensated measure \hat{N} is important since we can give meaning to the integral

$$\int_A f(x)\hat{N}(t, \mathbf{d}x), \quad (3.28)$$

for $A \subset E$ (not necessarily bounded below) by approximating it with the L^2 limit of (3.27) over a sequence of sets A_n which are bounded below. It is essential that we have the compensation term in (3.27) since the first integral in (3.27) may not converge if A is not bounded below.

The Lévy-Ito decomposition

If X is a Lévy process with triplet (b, A, ν) then there exists a standard Brownian motion W and an independent Poisson random measure, N , with intensity $\mathbf{d}s \times \nu(\mathbf{d}z)$ such that

$$X(t) = bt + \Sigma W(t) + \int_{|x| < 1} x \hat{N}(t, \mathbf{d}x) + \int_{|x| \geq 1} x N(t, \mathbf{d}x), \quad (3.29)$$

where $\Sigma\Sigma^\top := A$. The decomposition (3.29) is called the **Lévy-Ito decomposition** [3]. While the Lévy-Khintchine representation characterizes the distribution of a Lévy process, the Lévy-Ito decomposition characterizes the structure of the paths of a Lévy process. Consequently, it gives us intuition as to how a Lévy process can be integrated against. Firstly, if one notes that

$$\Sigma W(t) + \int_{|x|<1} x\hat{N}(t, dx) \quad (3.30)$$

is a martingale and

$$bt + \int_{|x|\geq 1} xN(t, dx) \quad (3.31)$$

is a finite variation process, it is clear to see that any Lévy process is a semimartingale, for which there exists the theory of stochastic integration [114]. Furthermore, we can define the stochastic integral term-wise from (3.29). For a nice enough stochastic process Y (e.g. predictable with certain integrability conditions [3]), we can define the integral of Y against the Lévy process X term-by-term from the decomposition (3.29):

$$\int_0^t Y(s)dX(s) := \int_0^t bY(s)ds + \int_0^t \Sigma Y(s)dW(s) \quad (3.32)$$

$$+ \int_0^t \int_{|x|<1} xY(s)\hat{N}(ds, dx) + \int_0^t \int_{|x|\geq 1} xY(s)N(ds, dx). \quad (3.33)$$

More generally, however, we can consider a so-called **Lévy-type stochastic integral** which has the form

$$\begin{aligned} \int_0^t G(s)ds + \int_0^t F(s)dW(s) + \int_0^t \int_{|x|<1} H(s, x)\hat{N}(dt, dx) \\ + \int_0^t \int_{|x|\geq 1} K(s, x)N(dt, dx), \end{aligned} \quad (3.34)$$

for suitable integrands G, F, H , and K . Analogously to Ito integrals, we can consider Lévy-type stochastic integrals as stochastic processes (in t) and we can consider the case when the integrands depend on the resulting stochastic process.

Consider the Lévy-driven SDE:

$$\begin{aligned}
X(t) = & X(0) + \int_0^t b(s, X(s-)) \mathrm{d}s + \int_0^t \sigma(s, X(s-)) \mathrm{d}W(s) \\
& + \int_0^t \int_{|x| < c} F(s, X(s-)) x \hat{N}(\mathrm{d}t, \mathrm{d}x) \\
& + \int_0^t \int_{|x| \geq c} G(s, X(s-)) x N(\mathrm{d}t, \mathrm{d}x), \tag{3.35}
\end{aligned}$$

where $c > 0$. The constant c is arbitrary; in the case that we have finite activity we can take $c = 0$. Existence and uniqueness depends on properties of the coefficients. We make the following two assumptions:

Assumption 3.1 (Lipschitz coefficients). *There exists a constant $L > 0$ such that for all $t \in [0, T]$ and all $x, y \in \mathbb{R}^d$,*

$$\begin{aligned}
& |b(t, x) - b(t, y)|^2 + |\sigma(t, x) - \sigma(t, y)|^2 \\
& + \int_{|z| < c} |F(t, x) - F(t, y)|^2 |z|^2 \nu(\mathrm{d}z) \leq L|x - y|^2. \tag{3.36}
\end{aligned}$$

Assumption 3.2 (Linear growth). *There exists a constant $C > 0$ such that for all $t \in [0, T]$ and all $x \in \mathbb{R}^d$,*

$$|b(t, x)|^2 + |\sigma(t, x)|^2 + \int_{|z| < c} |F(t, x)|^2 |z|^2 \nu(\mathrm{d}z) \leq C(1 + |x|^2). \tag{3.37}$$

Theorem 3.1. *Let W be a Wiener process on (Ω, \mathcal{F}, P) and let N be an independent Poisson random measure. Further, assume the functions b, σ, F and G satisfy Assumptions 3.1 and 3.2. Then there exists a unique càdlàg process X which satisfies (3.35).*

A proof of Theorem 3.1 can be found in [3]. Generalizations can be found in [3, 114].

Partial-integro differential equations

We now recall the connection between Lévy-driven SDEs and partial-integro differential equations (PIDEs) [3, Sec. 6.7].

Fix $0 \leq t_0 < T < \infty$ and for any $t \in [t_0, T]$ and any $x \in \mathbb{R}^d$, let $X = X_{t,x}$ be the unique solution of

$$dX = b(s, X)ds + \sigma(s, X)dW(s) + \int_{\mathbb{R}^q} F(t, X(s-))z\hat{N}(ds, dz), \quad (3.38)$$

with the initial condition $X_{t,x}(t) = x \in \mathbb{R}^d$. Under Assumptions 3.1 and 3.2 and for a regular enough test function, f , consider the following expectation as a function of the initial condition:

$$u(t, x) = \mathbb{E}[f(X_{t,x}(T))]. \quad (3.39)$$

The function u satisfies the (parabolic) partial-integro differential equation (PIDE):

$$\frac{\partial u}{\partial t} + L_t u = 0, \quad (t, x) \in [t_0, T) \times \mathbb{R}^d \quad (3.40)$$

$$u(T, x) = f(x), \quad x \in \mathbb{R}^d, \quad (3.41)$$

where L_t is the integro-differential operator

$$\begin{aligned} L_t u(t, x) := & \frac{1}{2} \text{tr}[\sigma^\top(t, x) \nabla^2 u(t, x) \sigma(t, x)] + \langle b(t, x), \nabla u(t, x) \rangle \\ & + \int_{\mathbb{R}^q} \left[u(t, x + F(t, x)z) - u(t, x) \right. \\ & \left. - \langle F(t, x)z, \nabla u(t, x) \rangle \mathbb{1}_{|z| \leq 1} \right] \nu(dz), \end{aligned} \quad (3.42)$$

see, for example [34, 3].

3.2.3 Artificial Neural Networks

In this section, we introduce some material on artificial neural networks. We follow the notation from [14].

For $L \in \mathbb{N}$, $N = (N_0, \dots, N_L) \in \mathbb{N}^{L+1}$ and a Lipschitz function $\rho :$

$\mathbb{R} \rightarrow \mathbb{R}$, a fully connected feed-forward neural network is defined by its architecture $a = (N, \rho)$ and its realization $\Phi_a : \mathbb{R}^{N_0} \times \mathbb{R}^{P(N)} \rightarrow \mathbb{R}^{N_L}$, where $P(N)$ is the number of parameters

$$P(N) := \sum_{l=1}^L N_l N_{l-1} + N_L. \quad (3.43)$$

The realization of the neural network architecture is of the form

$$\Phi_{a,\theta}(x) = W^{(L)} \left(\dots A^{(2)} \left(W^{(2)} \left(A^{(1)} \left(W^{(1)} x + b^{(1)} \right) \right) + b^{(2)} \right) \dots \right) + b^{(L)}, \quad (3.44)$$

where

$$W^{(l)} \in \mathbb{R}^{N_l \times N_{l-1}}, \quad b^{(l)} \in \mathbb{R}^{N_l}, \quad l = 1, \dots, L, \quad (3.45)$$

and $A^{(l)} : \mathbb{R}^{N_l} \rightarrow \mathbb{R}^{N_l}$ is defined by

$$A^{(l)}(x) = \left(\rho(x_1), \dots, \rho(x_{N_l}) \right)^\top, \quad l = 1, \dots, L-1, \quad (3.46)$$

and $\theta = ((W^{(l)}, b^{(l)}))_{l=1}^L \in \mathbb{R}^{P(N)}$. The neural network is said to have $(L-1)$ hidden layers. As the activation function ρ we use ReLU. For further details, see e.g. [59].

3.3 Variance reduction for stochastic differential equations

In the case of SDEs driven only by Brownian motion, the two variance reduction methods of importance sampling and control variates are well-known and have been studied in [105, 98, 101] and references therein.

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t_0 \leq t \leq T}, P)$ be a filtered probability space on which a d -dimensional Brownian motion $W(t)$ is defined. For $s \in [t_0, T]$ and $x \in \mathbb{R}^d$, consider the stochastic processes $X_{s,x}(t), Y_{s,x}(t), Z_{s,x}(t), t \geq s$, defined as

the solution to the system:

$$dX(t) = b(t, X)dt + \sigma(t, X)dW(t), \quad X(s) = x, \quad (3.47)$$

$$dY(t) = c(t, X)Y(t)dt, \quad Y(s) = 1, \quad (3.48)$$

$$dZ(t) = g(t, X)Y(t)dt, \quad Z(s) = 0, \quad (3.49)$$

where $b(t, x)$ is a d -dimensional vector, $\sigma(t, x)$ is a $d \times d$ matrix, and $c(t, x)$ and $g(t, x)$ are scalar functions, all with appropriate regularity properties [46, 48]. We are interested in calculating the expectation

$$u(s, x) := \mathbb{E}[f(X_{s,x}(T))Y_{s,x}(T) + Z_{s,x}(T)]. \quad (3.50)$$

Let us briefly justify our interest. As a concrete example of an application, consider the option pricing problem. In this example, the process X would be a model for the random path of a stock price (or any other market observable), Y would be a discounting factor, and f the payoff of an option. The process Z would represent the cumulative (discounted) running costs of holding the replicating portfolio. These running costs could be financing/storage costs or dividend payments. The expectation in (3.50) would then be the fair price of the option, under the prescribed financial model, determined by the coefficients in (3.47)-(3.49). More generally, however, the function $u : [t_0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is known [42, 46, 48] to satisfy the related Cauchy problem for the following parabolic PDE:

$$\frac{\partial u}{\partial t} + Lu + c(t, x)u + g(t, x) = 0, \quad (t, x) \in [t_0, T] \times \mathbb{R}^d \quad (3.51)$$

$$u(T, x) = f(x), \quad x \in \mathbb{R}^d, \quad (3.52)$$

where L is a differential operator of the form

$$Lu(t, x) := \frac{1}{2} \text{tr}[\sigma^\top(t, x) \nabla^2 u(t, x) \sigma(t, x)] + \langle b(t, x), \nabla u(t, x) \rangle. \quad (3.53)$$

Hence, if we wish to solve a problem of the form (3.51)-(3.52), we can appeal to its stochastic representation (3.47)-(3.50) and compute the expec-

tation in (3.50). This is a particularly effective strategy in high dimensions. Thus, the interest of computing (3.50) is clearly established.

To ease notation, let us denote the random variable whose expectation we wish to compute as

$$\Gamma := \Gamma_{s,x} := f(X_{s,x}(T))Y_{s,x}(T) + Z_{s,x}(T), \quad (3.54)$$

and, more generally, let us define the process

$$\Gamma(t) := \Gamma_{s,x}(t) := u(t, X_{s,x}(t))Y_{s,x}(t) + Z_{s,x}(t), \quad t \in [s, T], \quad (3.55)$$

noting that $\Gamma_{s,x}(T) = \Gamma_{s,x}$ and $\Gamma_{s,x}(s) = u(s, x)$. The process $\Gamma(t)$ will be a central object in the proofs of this section.

3.3.1 Control Variates

Note that for a suitable integrand G (say Lipschitz with linear growth), the Ito integral

$$\int_s^T G^\top(t, X(t))dW(t) \quad (3.56)$$

has zero expectation, hence it can serve as a control variate for the random variable Γ , whose expectation we would like to compute. Moreover, as we will show, G can be chosen in such a way so that the integral (3.56) is perfectly anti-correlated with Γ , which means that

$$\Gamma + \int_s^T G^\top(t, X(t))dW(t) \quad (3.57)$$

has zero variance (cf. equation (3.10)). We can incorporate this control variate into the SDE for Z :

$$dZ(t) = g(t, X)Y(t)dt + G^\top(t, X)dW(t), \quad Z(s) = 0, \quad (3.58)$$

with Γ defined as before but using (3.58). The method of control variates for SDEs was first considered in [105, 104]. Note the connection with the Martingale Representation theorem [108, 73] which ensures the existence

of an adapted process $\phi(t)$ such that

$$\Gamma = \mathbb{E}[\Gamma] + \int_s^T \phi^\top(t) dW(t). \quad (3.59)$$

3.3.2 Importance sampling

For a function $\mu : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, which satisfies the Novikov condition

$$\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^T |\mu(s, X(s))|^2 ds \right) \right] < \infty, \quad (3.60)$$

the process

$$\tilde{W}(t) = W(t) - \int_0^t \mu(s, X(s)) ds, \quad 0 \leq t \leq T, \quad (3.61)$$

is a Wiener process on $[0, T]$ under the measure \tilde{P} defined through the Radon-Nikodym derivative

$$\frac{d\tilde{P}}{dP} = \exp \left\{ \int_0^t \mu^\top(s, X(s)) dW(s) - \frac{1}{2} \int_0^t |\mu(s, X(s))|^2 ds \right\}, \quad (3.62)$$

see, for example, [73]. Hence, we can change the measure of the Wiener process driving X , and then compensate for this by multiplying by the Radon-Nikodym derivative (3.62), keeping $\mathbb{E}\Gamma$ fixed. This amounts to considering the system

$$dX = b(t, X)dt - \sigma(t, X)\mu(t, X)dt + \sigma(t, X)dW(t), \quad X(s) = x, \quad (3.63)$$

$$dY = c(t, X)Ydt + \mu^\top(t, X)YdW(t), \quad Y(s) = 1, \quad (3.64)$$

with Γ defined as before through (3.63), (3.64) and (3.49). The function μ completely determines the change of measure and, as was discussed in Section 3.2.1, should be chosen such that Γ has small variance. The method of importance sampling for SDEs was first considered in [56, 142].

3.3.3 Combining method

We now consider the combining method, i.e. using both control variates and importance sampling, and present the theoretical results concerning optimal choices of control variates and changes of measures that were alluded to in the previous sections. This method was first considered in [98] (see also [101]) Instead of the system (3.47)-(3.49), consider the system

$$dX = b(t, X)dt - \sigma(t, X)\mu(t, X)dt + \sigma(t, X)dW(t), \quad X(s) = x, \quad (3.65)$$

$$dY = c(t, X)Ydt + \mu^\top(t, X)YdW(t), \quad Y(s) = 1, \quad (3.66)$$

$$dZ = g(t, X)Ydt + G^\top(t, X)YdW(t), \quad Z(s) = 0, \quad (3.67)$$

where μ and G are d -dimensional vector functions with good analytical properties. The role of μ and G should be clear from the discussion in Sections 3.3.1 and 3.3.2. Note that the original system (3.47)-(3.49) is a special case of this system with $\mu = G = 0$. Of course, the case when only $\mu = 0$ corresponds to the method of control variates and the case when only $G = 0$ corresponds to the method of importance sampling (cf. Sections 3.3.1 and 3.3.2).

It is straightforward to show that while $\mathbb{E}\Gamma$ does not depend on μ or G , the variance of Γ does indeed depend on them. Moreover, the method can be made optimal, in the sense that there exists an optimal choice of μ and G , such that $\text{Var}\Gamma = 0$. This optimal choice, however, depends on (and consequently requires knowledge of) the full solution of the related PDE problem (3.51)-(3.52) and its spatial derivatives. This is summarized by the following theorem:

Theorem 3.2 (Milstein & Schoenmakers [98]). *If μ and G are such that*

$$u(t, x)\mu(t, x) + G(t, x) = -\sigma^\top(t, x)\nabla u(t, x) \quad (3.68)$$

for all $(t, x) \in [s, T] \times \mathbb{R}^d$, where u is the solution to the PDE (3.51)-(3.52), then $\text{Var}\Gamma_{s,x}(T) = 0$.

This theorem demonstrates a general possibility of perfect variance reduction and serves as a guidance towards choosing (some suboptimal but practical) μ or/and G .

Proof of Theorem 3.2. Applying Ito's lemma to $\Gamma(t)$, we get

$$\begin{aligned} d\Gamma(t) &= Y \left(\frac{\partial u}{\partial t} + \frac{1}{2} \text{tr}[\sigma \sigma^\top \nabla^2 u] + \langle b, \nabla u \rangle - \langle \sigma \mu, \nabla u \rangle \right) dt \\ &\quad + Y (\sigma^\top \nabla u)^\top dW + ucY dt + Y \langle \sigma \mu, \nabla u \rangle dt + Yu\mu^\top dW \\ &\quad + gY dt + G^\top Y dW, \end{aligned} \quad (3.69)$$

where the coefficients are evaluated at $(t, X(t))$. Noticing that the function u satisfies the PDE (3.51), we are left with

$$d\Gamma(t) = Y (\sigma^\top \nabla u + u\mu + G)^\top dW. \quad (3.70)$$

Writing (3.70) in integral form over $[s, T]$ and noting that $\Gamma(s) = u(s, x)$, we obtain

$$\Gamma(T) = u(s, x) + \int_s^T Y (\sigma^\top \nabla u + u\mu + G)^\top dW. \quad (3.71)$$

Then, by the Ito isometry,

$$\text{Var}\Gamma = \mathbb{E} \int_s^T Y^2(t) \|\sigma^\top \nabla u + u\mu + G\|^2 dt, \quad (3.72)$$

and if the condition (3.68) is satisfied, $\text{Var}\Gamma = 0$. Finally, note that taking expectations of (3.71) yields

$$\mathbb{E}\Gamma = u(s, x), \quad (3.73)$$

demonstrating that μ and G do not change the expected value of Γ . \square

Remark 3.1. *It is clear from Theorem 3.2 that we do not require both μ and G to be non-zero to achieve perfect variance reduction. In other words, the methods of importance sampling and control variates are both sufficient on their own to reduce the variance of Γ to zero. For example, for importance sampling we can*

take μ to satisfy

$$u(t, x)\mu(t, x) = -\sigma^\top(t, x)\nabla u(t, x),$$

and for control variates, we can take

$$G(t, x) = -\sigma^\top(t, x)\nabla u(t, x).$$

3.3.4 Numerical algorithm

We now discuss how to leverage the result of Theorem 3.2 to positive practical effect. We consider the (neural) SDEs

$$dX = b(t, X)dt + \sigma(t, X)dW(t), \quad X(s) = x, \quad (3.74)$$

$$dY = c(t, X)Ydt, \quad Y(s) = 1, \quad (3.75)$$

$$dZ = g(t, X)Ydt + G_\theta^\top(t, X)YdW(t), \quad Z(s) = 0, \quad (3.76)$$

where $G_\theta : [t_0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a neural network parameterization of G (see Section 3.2.3).

Provided that the system (3.47)-(3.49) has a unique strong solution, it is sufficient that G_θ be Lipschitz in x for the existence and uniqueness of a solution to (3.74)-(3.76). In the case of feed-forward architectures, which we will make use of, this amounts to G_θ having a Lipschitz activation function (see Section 3.2.3).

Comparably to before, define

$$\Gamma_\theta := f(X_{s,x}(T))Y_{s,x}(T) + Z_{s,x}(T). \quad (3.77)$$

The subscript θ in Γ_θ denotes the implicit dependence of Γ_θ on the parameters of the neural network, θ .

For any choice of θ , the expectation $\mathbb{E}\Gamma_\theta$, remains unchanged. Moreover, since Theorem 3.2 implies that there exists an optimal choice of G such that we have zero variance of Γ , our objective is to find parameters θ^* such that G_{θ^*} approximately satisfies (3.68) and hence $\text{Var}\Gamma_{\theta^*} \approx 0$. We will denote the true optimal value of G by G^* . In other words, from

Theorem 3.2, G^* satisfies

$$G^*(t, x) = -\sigma^\top(t, x)\nabla u(t, x). \quad (3.78)$$

Remark 3.2. *Let us briefly justify the choice of feed-forward neural networks as approximators. There exists many theoretical results on the expressivity of neural networks [35, 68, 7]. In particular, the universal approximation theorem given by Hornik et al. [68] states that for any finite measure on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$, the class of feed-forward neural networks, mapping from \mathbb{R}^d to \mathbb{R} with continuous and non-constant activation is dense in L^p . The function we would like to approximate, G^* , maps into \mathbb{R}^d but the above approximation theorem can be applied componentwise. Furthermore, there is a growing body of work on the numerical solution to PDEs using neural networks [124, 63, 52], in particular focused on the effectiveness of neural networks in solving high-dimensional problems, see for example [13, 58] and the references mentioned therein. It is precisely in these high-dimensional situations where one may resort to Monte Carlo methods, and the above cited research in deep learning provides a justification for making use of neural SDEs within this context.*

The problem of finding optimal parameters θ^* amounts to solving the optimization problem

$$\theta^* \in \arg \min_{\theta} \text{Var} \Gamma_{\theta}. \quad (3.79)$$

Of course, in any non-trivial situation $\text{Var} \Gamma_{\theta}$ cannot be evaluated analytically, nor can the random variable Γ_{θ} be simulated exactly. Instead we must use the empirical (sample) variance of independent realizations of an approximate random variable, $\bar{\Gamma}_{\theta}$, which is close to Γ_{θ} in the weak sense. Fixing a large $M_r \in \mathbb{N}$, the problem becomes

$$\theta^* \in \arg \min_{\theta} \text{Var}_{M_r} \bar{\Gamma}_{\theta}, \quad (3.80)$$

where $\text{Var}_{M_r}(\cdot)$ denotes the empirical variance over M_r realizations. The random variables $(\bar{\Gamma}_{\theta, m})_{m=1}^{M_r}$ are obtained by numerical integration of the

neural SDE system (3.74)-(3.76) together with (3.77). This optimization problem can be solved using a stochastic gradient-based algorithm (e.g. stochastic gradient descent), noting that the loss function

$$\mathcal{L}(\theta) := \text{Var}_{M_r} \bar{\Gamma}_\theta \quad (3.81)$$

is differentiable.

Once the parameters θ^* have been found, realizations of $\bar{\Gamma}_{\theta^*}$ can be simulated using a (potentially different) numerical integration scheme. The new MC estimator is then given by

$$\bar{u}(s, x) = M^{-1} \sum_{m=1}^M \bar{\Gamma}_{\theta^*, m}, \quad (3.82)$$

where $(\bar{\Gamma}_{\theta^*, m})_{m=1}^M$ are independent copies of $\bar{\Gamma}_{\theta^*}$.

This motivates a two-pass algorithm, where the first pass finds optimal parameters θ^* and the second pass uses these parameters to simulate low-variance random variables for the MC estimator (cf. the two-run algorithm in [100]). This is described in Algorithm 1. In the first pass, we simulate (using a numerical integration scheme) trajectories of the solution to the system (3.74)-(3.76) with $G_\theta = 0$ and store them in memory along with the random variables used in the scheme. These trajectories can then be used to simulate the θ -dependent term, $\int_s^T G_\theta(t, X) Y(t) dW(t)$, at each iteration of the stochastic optimization algorithm. This eliminates the need to simulate the entire system for each iteration of the stochastic optimization algorithm.

In the second pass, we simply simulate solutions to the system (3.74)-(3.76) using G_{θ^*} . Note that the numerical scheme used in the second pass does not need to have anything in common with the numerical scheme of the first pass. In particular, we find it to be effective to use a coarse grid (larger h) in the first pass, followed by a fine grid (smaller h) in the second pass (cf. a similar observation in [100] in the case of the linear regression-based algorithm).

Letting $V(t) = (X(t)^\top, Y(t), Z(t))^\top$ and $v = (x^\top, 1, 0)^\top$, we consider numerical methods with a uniform step-size $h > 0$ of the form:

$$\bar{V}_{k+1} = \bar{V}_k + A(t_k, \bar{V}_k, h, \xi_k), \quad (3.83)$$

$$\bar{V}_0 = V(s) = v, \quad (3.84)$$

for some Borel measurable, vector-valued function A and random variables $(\xi_n)_{n \geq 0}$, where ξ_0 is independent of \bar{V}_0 and ξ_k is independent of $(\xi_n)_{0 \leq n < k}$ and $(\bar{V}_n)_{0 \leq n \leq k}$. We denote the scheme by (A, h) .

Algorithm 1: Neural control variate method for Brownian-driven SDEs

Result: MC approximation of the solution to the PDE (3.51)-(3.52)

at the point (s, x_0) , $\bar{u}(s, x_0)$

Initialize: Number of trials for first-pass M_r , number for trials for second-pass M , numerical scheme for first-pass (A_r, h_r) , numerical scheme for second-pass (A, h)

for $m \leftarrow 1$ **to** M_r **do**

Initialize: $\bar{X}_0, \bar{Y}_0, \bar{Z}_0 \leftarrow x_0, 1, 0$

Compute: $(t_k, \bar{X}_k, \bar{Y}_k, \bar{Z}_k)_{0 < k < N}^m$ and $(\bar{\Gamma})^m$ according to (3.74)-(3.76) with $G_\theta = 0$ and the scheme (A_r, h_r)

Store: $(t_k, \bar{X}_k, \bar{Y}_k, \bar{Z}_k)_{0 < k < N}^m$ and $(\bar{\Gamma})^m$ and random variables $(\xi_k)_{0 \leq k < N}$

end

Compute: $\theta^* = \arg \min_{\theta \in \Theta} \text{Var}_{M_r} \bar{\Gamma}_\theta$ by the stochastic optimization algorithm using the stored trajectories and random variables to compute $\bar{\Gamma}_\theta$ with the scheme (A_r, h_r)

for $m \leftarrow 1$ **to** M **do**

Initialize: $\bar{X}_0, \bar{Y}_0, \bar{Z}_0 \leftarrow x_0, 1, 0$

Compute: $\bar{\Gamma}_{\theta^*, m}$ using (3.74)-(3.76) with G_{θ^*} and the numerical scheme (A, h)

Store: Updated sample statistics of $\bar{\Gamma}_{\theta^*}$

end

Return: $\bar{u}(s, x_0) = M^{-1} \sum_{m=1}^M \bar{\Gamma}_{\theta^*, m}$

Remark 3.3. *For the numerical methods, Algorithm 1 and later Algorithm 2, we only consider the case of control variates ($\mu = 0$). The principal reason for this is computational efficiency. What makes Algorithm 1 effective is the fact that a large batch of trajectories of the SDEs can be simulated once, in parallel, before training. However, in the importance sampling case, trajectories of X, Y and Z depend on the neural network parameters, θ . Thus, each time θ is updated new trajectories have to be sampled, which makes the approach computationally intractable.*

In practice the learned G_{θ^*} is not equal to the optimal G^* satisfying (3.78) due to errors of deep learning (finite size of the network, limited training set and accuracy of the stochastic optimization algorithm) and due to the numerical integration error. Recall from (3.72):

$$\text{Var}\Gamma_{\theta^*} = \mathbb{E} \int_s^T Y_{s,x_0}^2(t) \|\sigma \nabla u + G_{\theta^*}\|^2 dt.$$

Then, using (3.78), we get that the standard deviation of Γ_{θ^*} gives the error of Γ_{θ^*} in a weighted norm:

$$\sqrt{\text{Var}\Gamma_{\theta^*}} = \left(\mathbb{E} \int_s^T Y_{s,x_0}^2(t) |G_{\theta^*} - G^*|^2 dt \right)^{1/2}.$$

Consequently, ignoring the error of numerical integration, we can view

$$\text{Err}_{G_{\theta^*}} = \frac{\sqrt{\text{Var}\Gamma_{\theta^*}}}{\mathbb{E}\Gamma_{\theta^*}} \tag{3.85}$$

as the appropriated relative error of the trained G_{θ^*} .

3.4 Lévy-driven stochastic differential equations

We now turn our attention to the case of SDEs driven by a more general Lévy noise, i.e., SDEs driven by both Wiener and Poisson processes. Such systems arise as modelling tools frequently in mathematical finance [34]. For example, the well known models of Merton [94] and Kou [82] fall into this category, among others (see [34] for a comprehensive overview).

In this case, we have the corresponding Cauchy problem for the PIDE:

$$\frac{\partial u}{\partial t} + Lu + c(t, x)u + g(t, x) = 0, \quad (t, x) \in [t_0, T] \times \mathbb{R}^d, \quad (3.86)$$

$$u(T, x) = f(x), \quad x \in \mathbb{R}^d, \quad (3.87)$$

where L is a partial integro-differential operator of the form

$$\begin{aligned} Lu(t, x) := & \frac{1}{2} \text{tr}[a(t, x) \nabla^2 u(t, x)] + \langle b(t, x), \nabla u(t, x) \rangle \\ & + \int_{\mathbb{R}^q} \left[u(t, x + F(t, x)z) - u(t, x) \right. \\ & \left. - \langle F(t, x)z, \nabla u(t, x) \rangle \mathbb{1}_{|z| \leq 1} \right] \nu(\mathbf{d}z). \end{aligned} \quad (3.88)$$

Here $a(t, x)$ is a symmetric positive semidefinite $d \times d$ -matrix, $b(t, x)$ is a d -dimensional vector, $c(t, x)$ and $g(t, x)$ are scalar functions, $F(t, x)$ is a $d \times q$ -matrix, and ν is a Lévy measure. We allow for the possibility that ν is of infinite intensity, i.e. we may have $\nu(B(0, r)) = \infty$ for some $r > 0$. Conditions for the existence and uniqueness of a solution to the problem (3.86)-(3.87) can be found in [51].

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t_0 \leq t \leq T}, P)$ be a filtered probability space on which a d -dimensional standard Wiener process $w(t)$ and a Poisson random measure N on $[0, \infty) \times \mathbb{R}^q$ with intensity $ds \times \nu(\mathbf{d}z)$ are defined. Let \hat{N} denote the corresponding Poisson random measure with compensated small jumps, see Section 3.2.2.

We assume that the problem (3.86)-(3.87) admits a classical solution, $u \in C^{1,2}([t_0, T] \times \mathbb{R}^d)$. It has the probabilistic representation (see e.g. [3, 34]) given by

$$u(s, x) = \mathbb{E}[f(X_{s,x}(T))Y_{s,x}(T) + Z_{s,x}(T)], \quad (3.89)$$

where $X_{s,x}(t), Y_{s,x}(t), Z_{s,x}(t), s \leq t \leq T$, is the solution of the system of

SDEs:

$$dX = b(t, X)dt + \sigma(t, X)dw(t) + \int_{\mathbb{R}^q} F(t, X(t-))z\hat{N}(dt, dz), \quad (3.90)$$

$$dY = c(t, X)Y(t-)dt, \quad (3.91)$$

$$dZ = g(t, X)Y(t-)dt, \quad (3.92)$$

with $X(s) = x$, $Y(s) = 1$, and $Z(s) = 0$. Here the matrix $\sigma(t, x)$ is a solution of the equation $a(t, x) = \sigma(t, x)\sigma^\top(t, x)$. Note that when we write e.g. $b(t, X)$ we now mean $b(t, X(t-))$, i.e. we are now using the left limits of the stochastic processes as the point at which we evaluate the coefficients. Recall from Section 3.2 that $X(t-) := \lim_{s \nearrow t} X(s)$.

In order to be able to efficiently simulate approximate realizations of the process X , we follow [4] (see also [81, 36]) and consider a modified process X^ε , where the small jumps of X are approximated with an additional diffusion component. The resulting X^ε is a jump-diffusion with only finitely many jumps on any finite time interval. Intuitively, we are replacing the cumulative effect of the small jumps (which is hard to simulate) with a Brownian motion with covariance structure that matches the small jumps we are removing. Since these small jumps are happening infinitely often and are small enough, this is justified by a functional central limit theorem, see [4]

Let $W(t)$ be a q -dimensional Brownian motion independent of w and N . Then the process $X_{s,x}^\varepsilon(t)$ and the corresponding $Y_{s,x}^\varepsilon(t)$, $Z_{s,x}^\varepsilon(t)$ are defined as the solution to

$$dX^\varepsilon = b(t, X^\varepsilon(t-)) - F(t, X^\varepsilon(t-))\gamma_\varepsilon dt + \sigma(t, X^\varepsilon(t-))dw(t) \\ + F(t, X^\varepsilon(t-))\beta_\varepsilon dW(t) + \int_{|z| \geq \varepsilon} F(t, X^\varepsilon(t-))zN(dt, dz), \quad (3.93)$$

$$dY^\varepsilon = c(t, X^\varepsilon)Y^\varepsilon(t-)dt, \quad (3.94)$$

$$dZ^\varepsilon = g(t, X^\varepsilon)Y^\varepsilon(t-)dt, \quad (3.95)$$

with the same initial conditions as in (3.90)-(3.92). The vector γ_ε is defined component-wise as

$$\gamma_\varepsilon^i = \int_{\varepsilon \leq |z| \leq 1} z^i \nu(dz), \quad (3.96)$$

and β_ε is defined by

$$B_\varepsilon^{ij} = \int_{|z| < \varepsilon} z^i z^j \nu(dz), \quad \beta_\varepsilon \beta_\varepsilon^\top = B_\varepsilon. \quad (3.97)$$

As before we introduce the random variable, whose expectation we would like to compute, as

$$\Gamma^\varepsilon := \Gamma_{s,x}^\varepsilon := f(X_{s,x}^\varepsilon(T))Y_{s,x}^\varepsilon(T) + Z_{s,x}^\varepsilon(T), \quad (3.98)$$

and, more generally,

$$\Gamma^\varepsilon(t) := \Gamma_{s,x}^\varepsilon(t) := u(t, X_{s,x}^\varepsilon(t))Y_{s,x}^\varepsilon(t) + Z_{s,x}^\varepsilon(t), \quad (3.99)$$

noting that $\Gamma^\varepsilon = \Gamma^\varepsilon(T)$. We note the abuse of notation here: Γ^ε refers to a random variable and $\Gamma^\varepsilon(\cdot)$ refers to a process.

We can approximate the solution, $u(t, x)$, to the PIDE (3.86)-(3.87) by

$$u(s, x) \approx u^\varepsilon(s, x) := \mathbb{E}[\Gamma^\varepsilon(T)]. \quad (3.100)$$

It is shown in [36] (see also [81]) that $u^\varepsilon(t, x)$ is a good approximation for $u(t, x)$, whose accuracy is controlled by ε :

$$\|u(s, x) - u^\varepsilon(s, x)\| \leq K \int_{|z| \leq \varepsilon} \|z\|^3 \nu(dz), \quad (s, x) \in [0, T] \times \mathbb{R}^d,$$

where $K > 0$ does not depend on ε . For instance, for a tempered α -stable process (see details in [36, Example 2.1]) this error is of order $O(\varepsilon^{3-\alpha})$, $\alpha \in (0, 2)$.

The PIDE problem for u^ε is given by

$$\frac{\partial u}{\partial t} + L_\varepsilon u^\varepsilon + c(t, x)u^\varepsilon + g(t, x) = 0, \quad (3.101)$$

$$u^\varepsilon(T, x) = f(x), \quad (3.102)$$

where L_ε is the partial integro-differential operator

$$\begin{aligned} L_\varepsilon v(t, x) := & \frac{1}{2} \text{tr} \left[(a(t, x) + F(t, x)B_\varepsilon F^\top(t, x)) \nabla^2 v(t, x) \right] \\ & + \langle b(t, x) - F(t, x)\gamma_\varepsilon, \nabla v(t, x) \rangle \\ & + \int_{|z| \geq \varepsilon} \left[v(t, x + F(t, x)z) - v(t, x) \right] \nu(\mathbf{d}z). \end{aligned} \quad (3.103)$$

In the same spirit as in the Brownian case considered in Section 3.3, we modify the system (3.93)-(3.95) to allow for variance reduction of Γ^ε without changing the expectation of Γ^ε . In this case we have three sources of noise, namely w , W , and N . For the Brownian motions, we use importance sampling and control variate analogously to the results of Section 3.3. In addition, we introduce a control variate to deal with the Poisson random measure N . To this end, we introduce the five auxiliary functions, $\mu_w : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\mu_W : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^q$, $G_w : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, $G_W : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^q$, and $G_N : [0, \infty) \times \mathbb{R}^{d+q} \rightarrow \mathbb{R}$, which can be arbitrary except for some regularity conditions.

Consider now the system

$$\begin{aligned} dX^\varepsilon = & \left[b(t, X^\varepsilon) - F(t, X^\varepsilon)\gamma_\varepsilon - \sigma(t, X^\varepsilon)\mu_w(t, X^\varepsilon) \right. \\ & \left. - F(t, X^\varepsilon)\beta_\varepsilon\mu_W(t, X^\varepsilon) \right] dt \end{aligned} \quad (3.104)$$

$$\begin{aligned} & + \sigma(t, X^\varepsilon)dw(t) + F(t, X^\varepsilon)\beta_\varepsilon dW(t) \\ & + \int_{|z|\geq\varepsilon} F(t, X^\varepsilon)zN(dt, dz), \end{aligned}$$

$$dY^\varepsilon = c(t, X^\varepsilon)Y^\varepsilon dt + Y^\varepsilon\mu_w^\top(t, X^\varepsilon)dw(t) + Y^\varepsilon\mu_W^\top(t, X^\varepsilon)dW(t) \quad (3.105)$$

$$\begin{aligned} dZ^\varepsilon = & g(t, X^\varepsilon)Y^\varepsilon dt + G_w^\top(t, X^\varepsilon)Y^\varepsilon dw(t) \\ & + G_W^\top(t, X^\varepsilon)Y^\varepsilon dW(t) + Y^\varepsilon \int_{|z|\geq\varepsilon} G_N(t, X^\varepsilon, z)N(dt, dz) \\ & - Y^\varepsilon \int_{|z|\geq\varepsilon} G_N(t, X^\varepsilon, z)\nu(dz)dt. \end{aligned} \quad (3.106)$$

Note that the previous system (3.93)-(3.95) is a special case of the above system when all five auxiliary functions are set to zero.

Our aim now is to replicate Theorem 3.2 in this setting, i.e. to show that these auxiliary functions can be chosen in such a way that the variance of $\Gamma_{s,x}^\varepsilon(T)$ becomes zero.

Theorem 3.3. *If the functions μ_w , μ_W , G_w , G_W , and G_N satisfy*

$$u^\varepsilon(t, x)\mu_w(t, x) + G_w(t, x) = -\sigma^\top(t, x)\nabla u^\varepsilon(t, x), \quad (3.107)$$

$$u^\varepsilon(t, x)\mu_W(t, x) + G_W(t, x) = -\beta_\varepsilon^\top F^\top(t, x)\nabla u^\varepsilon(t, x), \quad (3.108)$$

$$G_N(t, x, z) = u^\varepsilon(t, x) - u^\varepsilon(t, x + F(t, x)z), \quad (3.109)$$

for all $(t, x, z) \in [s, T] \times \mathbb{R}^d \times \mathbb{R}^q$, then

$$\text{Var}[\Gamma_{s,x}^\varepsilon(T)] = 0. \quad (3.110)$$

Moreover,

$$u^\varepsilon(s, x) = \Gamma_{s,x}^\varepsilon(T). \quad (3.111)$$

Proof. Applying Ito's formula to $\Gamma_{s,x}^\varepsilon(t)$, we obtain

$$\begin{aligned}
\Gamma_{s,x}^\varepsilon(T) &= u^\varepsilon(s, x) \\
&+ \int_s^T Y^\varepsilon \left[\frac{\partial u^\varepsilon}{\partial t} + \frac{1}{2} \sum_{i,j=1}^d a^{ij} \frac{\partial^2 u^\varepsilon}{\partial x^i \partial x^j} + \langle b, \nabla u^\varepsilon \rangle \right. \\
&\quad \left. - \langle F \gamma_\varepsilon, \nabla u^\varepsilon \rangle + cu^\varepsilon + g \right] dt \\
&+ \frac{1}{2} \int_s^T Y^\varepsilon \sum_{i,j=1}^d (F B_\varepsilon F^\top)^{ij} \frac{\partial^2 u^\varepsilon}{\partial x^i \partial x^j} dt \\
&+ \int_s^T Y^\varepsilon (\nabla^\top u^\varepsilon \sigma + u^\varepsilon \mu_w^\top) dw(t) \\
&+ \int_s^T Y^\varepsilon (\nabla u^{\varepsilon \top} F \beta_\varepsilon + u^\varepsilon \mu_w^\top) dW(t) + \int_s^T Y^\varepsilon G_w^\top dw(t) \\
&+ \int_s^T Y^\varepsilon G_W^\top dW(t) \\
&+ \int_s^T \int_{|z| \geq \varepsilon} Y^\varepsilon [u^\varepsilon(t, X^\varepsilon(t-) + Fz) - u^\varepsilon] N(dt, dz) \\
&+ \int_s^T \int_{|z| \geq \varepsilon} Y^\varepsilon G_N N(dt, dz) - \int_s^T \int_{|z| \geq \varepsilon} Y^\varepsilon G_N \nu(dz) dt.
\end{aligned}$$

Since u^ε satisfies the PIDE (3.101)-(3.102), we arrive at

$$\begin{aligned}
\Gamma_{s,x}^\varepsilon(T) &= u^\varepsilon(s, x) \tag{3.112} \\
&+ \int_s^T \int_{|z| \geq \varepsilon} Y^\varepsilon [u^\varepsilon(t, X^\varepsilon + Fz) - u^\varepsilon] (N(dt, dz) - \nu(dz) dt) \\
&+ \int_s^T \int_{|z| \geq \varepsilon} Y^\varepsilon G_N (N(dt, dz) - \nu(dz) dt) \\
&+ \int_s^T Y^\varepsilon (\nabla^\top u^\varepsilon \sigma + G_w^\top + u^\varepsilon \mu_w^\top) dw(t) \\
&+ \int_s^T Y^\varepsilon (\nabla u^{\varepsilon \top} F \beta_\varepsilon + G_W^\top + u^\varepsilon \mu_W^\top) dW(t).
\end{aligned}$$

Then, it is not difficult to see that

$$\mathbb{E} \Gamma_{s,x}^\varepsilon(T) = u^\varepsilon(s, x), \tag{3.113}$$

regardless of the choice of the auxiliary functions. Moreover, if the conditions (3.107)-(3.109) are satisfied, the integrands in (3.112) become zero

and

$$\text{Var } \Gamma_{s,x}^\varepsilon(T) = 0. \quad (3.114)$$

□

3.4.1 Numerical algorithm

Theorem 3.3 can be used in the same way as Theorem 3.2 was used to motivate Algorithm 1. Theorem 3.3 shows that via an optimal choice of G_w , G_W , and G_N (here we take $\mu_w = \mu_W = 0$; see Remark 3.3) we can achieve a system of SDEs such that the random variable of interest, Γ^ε , has zero variance. Moreover, no matter the choice of G_w , G_W , or G_N we do not introduce any additional bias into the MC approximation which is clear from (3.113). We propose an algorithm whose objective is to find a good approximation of an optimal choice of the functions G_w , G_W , and G_N and which then uses this approximation to simulate low-variance realizations of Γ^ε .

As before, we parameterize G_w , G_W , and G_N with feed-forward neural networks $G_{w,\theta}$, $G_{W,\theta}$, and $G_{N,\theta}$, respectively. We obtain the neural SDEs

$$\begin{aligned} dX^\varepsilon &= b(t, X^\varepsilon) - F(t, X^\varepsilon)\gamma_\varepsilon dt + \sigma(t, X^\varepsilon)dw(t) \\ &\quad + F(t, X^\varepsilon)\beta_\varepsilon dW(t) + \int_{|z|\geq\varepsilon} F(t, X^\varepsilon)zN(dt, ds), \end{aligned} \quad (3.115)$$

$$dY^\varepsilon = c(t, X^\varepsilon)Y^\varepsilon dt, \quad (3.116)$$

$$\begin{aligned} dZ^\varepsilon &= g(t, X^\varepsilon)Y^\varepsilon dt + G_{w,\theta}^\top(t, X^\varepsilon)Y^\varepsilon dw(t) \\ &\quad + G_{W,\theta}^\top(t, X^\varepsilon)Y^\varepsilon dW(t) + Y^\varepsilon \int_{|z|\geq\varepsilon} G_{N,\theta}(t, X^\varepsilon, z)N(dt, dz) \\ &\quad - Y^\varepsilon \int_{|z|\geq\varepsilon} G_{N,\theta}(t, X^\varepsilon, z)\nu(dz)dt. \end{aligned} \quad (3.117)$$

Define

$$\Gamma_\theta^\varepsilon = f(X^\varepsilon(T))Y^\varepsilon(T) + Z^\varepsilon(T). \quad (3.118)$$

Letting $V(t) = (X(t)^\top, Y(t), Z(t))^\top$ and $v = (x^\top, 1, 0)^\top$, we consider

a numerical scheme of the form

$$\bar{V}_{k+1} = \bar{V}_k + A(t_k, \bar{V}_k, h, \xi_k), \quad (3.119)$$

$$\bar{V}_0 = V(s) = v, \quad (3.120)$$

where a deterministic $h > 0$ is a maximum step size, A is a Borel measurable function, and $(\xi_n)_{n \geq 0}$ are appropriately chosen random variables taking into account randomness coming from the Wiener and Poisson processes. We allow for the case of a restricted jump-adapted numerical scheme as e.g. in [36] when, in particular, the number of time steps is random (see further details in Section 3.5.3).

Algorithm 2 for the Lévy noise case is analogous to the Brownian motion case (Algorithm 1), except we have more sources of noise.

Algorithm 2: Neural control variate method for Lévy-driven SDEs

Result: MC approximation of the solution $u(s, x_0)$ to the PIDE (3.86)-(3.87)

Initialize: Number of trials for first-pass M_r , number for trials for second-pass M , numerical scheme for first-pass (A_r, h_r) , numerical scheme for second-pass (A, h)

for $m \leftarrow 1$ **to** M_r **do**

Initialize: $\bar{X}_0^\varepsilon, \bar{Y}_0^\varepsilon, \bar{Z}_0^\varepsilon \leftarrow x_0, 1, 0$

Compute: $(t_k, \bar{X}_k^\varepsilon, \bar{Y}_k^\varepsilon, \bar{Z}_k^\varepsilon)_{0 \leq k < N}^m$ and $(\bar{\Gamma}^\varepsilon)^m$ according to (3.115)-(3.117) with $G_w = G_W = G_N = 0$ and the scheme (A_r, h_r)

Store: $(t_k, \bar{X}_k^\varepsilon, \bar{Y}_k^\varepsilon, \bar{Z}_k^\varepsilon)_{0 \leq k < N}^m$ and $(\bar{\Gamma}^\varepsilon)^m$ and random variables $(\xi_k)_{0 \leq k < N}$

end

Compute: $\theta^* = \arg \min_{\theta \in \Theta} \text{Var}_{M_r} \bar{\Gamma}_\theta^\varepsilon$ by the stochastic optimization algorithm using the stored trajectories and random variables to compute $\bar{\Gamma}_\theta^\varepsilon$ with the scheme (A_r, h_r) .

for $m \leftarrow 1$ **to** M **do**

Initialize: $\bar{X}_0^\varepsilon, \bar{Y}_0^\varepsilon, \bar{Z}_0^\varepsilon \leftarrow x_0, 1, 0$

Compute: $\bar{\Gamma}_{\theta^*, m}^\varepsilon$ using (3.115)-(3.117) with $G_w, \theta^*, G_W, \theta^*, G_N, \theta^*$ and the numerical scheme (A, h)

Store: Updated sample statistics of $\bar{\Gamma}_{\theta^*}^\varepsilon$

end

Return: $\bar{u}^\varepsilon(s, x_0) = M^{-1} \sum_{m=1}^M \bar{\Gamma}_{\theta^*, m}^\varepsilon$

We note that Remark 3.3 is also applicable in the case of Algorithm 2

Remark 3.4. When implementing Algorithm 2, it is computationally expensive to evaluate the double integrals in (3.117). We found that, in practice, replacing the function $G_{N,\theta} : \mathbb{R}^{d+q+1} \rightarrow \mathbb{R}$ with a linear approximation of $G_{N,\theta}$ in z , $g_{N,\theta} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^q$ is more efficient. In other words,

$$G_{N,\theta}(t, x, z) \approx g_{N,\theta}(t, x)^\top z. \quad (3.121)$$

This means that the inner integrals need not be numerically approximated at each time step since the values can be computed in advance (and in many cases will be known analytically). Explicitly, the integral in (3.117) becomes

$$\int_{|z| \geq \varepsilon} G_{N,\theta}(t, x, z) \nu(\mathbf{d}z) \approx g_{N,\theta}(t, x)^\top \int_{|z| \geq \varepsilon} z \nu(\mathbf{d}z), \quad (3.122)$$

and $\int_{|z| \geq \varepsilon} z \nu(\mathbf{d}z)$ does not have to be computed at each time step.

Analogously to the discussion on the relative error in the diffusion case at the end of Section 3.3, let us consider the error of the trained G_{w,θ^*} , G_{W,θ^*} , G_{N,θ^*} in comparison with the optimal G_w^* , G_W^* , G_N^* satisfying the conditions (3.107)-(3.109). It is not difficult to obtain using (3.112) and Ito's formula that

$$\begin{aligned} \text{Var} \Gamma_{\theta^*}^\varepsilon &= E \int_s^T (Y_{s,x_0}^\varepsilon(t))^2 \left[(\nabla^\top u^\varepsilon \sigma + G_{w,\theta^*}^\top)^2 + (\nabla u^{\varepsilon \top} F \beta_\varepsilon + G_{W,\theta^*}^\top)^2 \right. \\ &\quad \left. + \int_{|z| \geq \varepsilon} (u^\varepsilon(t, X^\varepsilon(t-) + Fz) - u^\varepsilon + G_{N,\theta^*})^2 \nu(\mathbf{d}z) \right] dt. \end{aligned}$$

Then, using (3.107)-(3.109), we get that the standard deviation of $\Gamma_{\theta^*}^\varepsilon$ gives the weighted error of $\Gamma_{\theta^*}^\varepsilon$:

$$\begin{aligned} \sqrt{\text{Var} \Gamma_{\theta^*}^\varepsilon} &= \left(E \int_s^T (Y_{s,x_0}^\varepsilon)^2(t) \left[\|G_{w,\theta^*} - G_w^*\|^2 + \|G_{W,\theta^*} - G_W^*\|^2 \right. \right. \\ &\quad \left. \left. + \int_{|z| \geq \varepsilon} (G_{N,\theta^*} - G_N^*)^2 \nu(\mathbf{d}z) \right] dt \right)^{1/2}. \end{aligned}$$

Consequently (analogously to (3.123)), ignoring the error of numerical

integration, we can view

$$\text{Err}_{G_{\theta^*}} = \frac{\sqrt{\text{Var}\Gamma_{\theta^*}^\varepsilon}}{\mathbb{E}\Gamma_{\theta^*}^\varepsilon} \quad (3.123)$$

as the appropriated relative error of the trained G_{w,θ^*} , G_{W,θ^*} and G_{N,θ^*} .

3.5 Numerical experiments

In this section, we provide several numerical examples from computational finance to demonstrate efficiency of Algorithms 1 and 2. We consider the cases of the SDEs driven by Lévy processes with infinite activity, finite activity, and being driven only by Brownian motion.

In the paper’s spirit of the ‘on-the-fly’ variance reduction algorithm, we propose to fix an ANN architecture prior to any knowledge of the particular problem. Of course, better results can be achieved if a particular architecture is chosen and/or tuned for each problem. We propose to use a fully-connected feed forward ANN with ReLU activation (see Section 3.2.3). The number of hidden layers and the size of the hidden layers are chosen via a hyperparameter search detailed in Section 3.5.1 and the values selected are shown in Table 1. In general, there is a trade-off between the neural network complexity and the time it takes to fit the model, which usually results in very deep neural networks being impractical for the algorithms of this paper. We also make use of batch normalization before the first layer of the network. For the optimization, we use the Adam algorithm [78] with a fixed learning rate of $\eta = 10^{-3}$. All experiments are performed on an NVidia Tesla V100 GPU.

For each experiment, we compare the neural control variate algorithm with vanilla MC. By vanilla MC, we mean MC without the use of control variates or any other variance reduction techniques. This comparison serves as a convenient benchmark since, as a rule, the additional training run of Algorithm 1/2 must be less computationally expensive than simply increasing the number of simulations of vanilla MC. The implementation is carried out in PyTorch and the code is available at

https://github.com/piers-hinds/sde_mc.

Both the MC simulations and the training of the ANNs are done on a GPU. MC is particularly efficient to implement in a GPU-supported scientific computing library like PyTorch since each MC simulation can be carried out independently, i.e. in parallel. The corresponding implementation is at https://github.com/piers-hinds/sde_mc (see the `SdeSolver` class which can solve multiple independent trajectories in parallel).

Jump-adapted schemes pose a unique challenge in order to simulate trajectories in parallel, since each trajectory may have a different number of time steps when a jump-adapted numerical scheme is used. Our solution involves storing the time points corresponding to each trajectory, as well as the trajectory itself. Full details of the implementation can be found at https://github.com/piers-hinds/sde_mc (see the `JumpDiffusionSolver` class).

In addition, in Section 3.5.5, we compare the neural control variate algorithm with Multilevel Monte Carlo method [55] as well as a crude control variate method [57].

In each experiment, the network is trained on $M_r = 3 \cdot 10^4$ trajectories for a maximum of 20 epochs before being used to generate trajectories. We propose and use a stopping rule for the training which is explained in Section 3.5.1; the rule allows for the termination of the training early if it becomes too slow. The training trajectories are chosen to have 5 times larger time-step sizes than the trajectories used in the final MC estimation, i.e. $h_r = 5h$, we refer to this value as the step factor. We use a batch size of 2000 in all experiments. We summarize the network hyperparameters in Table 1. In each experiment, the reported time taken includes training time.

3.5.1 Discussion of deep learning setup

In this section, we provide details concerning hyperparameter estimation and a stopping rule to terminate the training procedure used in the

Table 1: Hyperparameters and their chosen values

Hyperparameter	Value
Number of hidden layers	3
Hidden layer size	50
Step factor	5
Training data size	$3 \cdot 10^4$

experiments of Section 3.5.

Fix some tolerance level, $\epsilon > 0$. For the $(1 - \alpha) \times 100\%$ confidence interval of the MC error, let $\mathcal{C} = \mathcal{C}(\epsilon, \alpha)$ denote the total cost of the control variate method (i.e., Algorithm 1/2) required to reach this tolerance level. The total cost \mathcal{C} has two parts: the cost of training $\mathcal{C}_{\text{train}}$ and the cost of running M independent trajectories to compute the quantity of interest $\mathbb{E}\Gamma$. The MC tolerance is equal to

$$\epsilon = \Phi^{-1}(1 - \alpha/2) \frac{\sqrt{\text{Var}\bar{\Gamma}_{\theta^*}}}{\sqrt{M}},$$

and the cost to achieve the corresponding MC simulation is M multiplied by the cost to run a single trajectory. Then, the total cost \mathcal{C} can be conveniently expressed as

$$\mathcal{C} = \mathcal{C}_{\text{train}} + \Phi^{-1}(1 - \alpha/2)^2 \mathcal{C}_{\text{batch}} \frac{\text{Var}\bar{\Gamma}_{\theta^*}}{\epsilon^2 S_{\text{batch}}}, \quad (3.124)$$

where $\mathcal{C}_{\text{batch}}$ is the cost to sample one batch and S_{batch} is the size of the batch (note that the cost of simulating a single trajectory is equivalent to $\mathcal{C}_{\text{batch}}/S_{\text{batch}}$).

Hyperparameter estimation

In order to choose hyperparameters, i.e. the size of hidden layers, size of training sample, etc., we require a method of evaluating the effectiveness of the variance reduction. Naturally, we would like to choose hyperparameters to minimize the cost \mathcal{C} . Our approach is to perform the hyperparameter optimization on some of the simple examples and

then use these hyperparameters on the remaining ‘unseen’ examples. We hope that this demonstrates that the algorithms proposed can be used in a black-box fashion. Of course, for optimal performance one could select hyperparameters for each individual example.

We identify the following hyperparameters: the number of hidden layers in each neural network and the size of these hidden layers (see Section 3.2.3); the reduction factor of the number of steps in the training data compared to the number of steps in the final simulations, which we call the step factor; and the size of the training data. Note that for the hidden layer size, we add d neurons to the number stated, where d is the dimension of X ; for example, when we choose a hidden layer size of 50 the neural network has $50 + d$ neurons in each layer. We list each hyperparameter and its corresponding search space in Table 2.

Table 2: Hyperparameters and their possible values.

Hyperparameter	Search space
Number of hidden layers	$\{2, 3, 4\}$
Hidden layer size	$\{20, 30, 40, 50, 60, 70\}$
Step factor	$\{5, 10, 15, 20, 25, 30\}$
Training data size	$\{10^4, 2 \cdot 10^4, 3 \cdot 10^4, 4 \cdot 10^4, 5 \cdot 10^4, 6 \cdot 10^4\}$

In order to determine suitable values of the hyperparameters we minimize the cost \mathcal{C} (with $\alpha = 0.05$, $\epsilon = 10^{-3}$) averaged over four test examples from Section 3.5. Specifically, we use the experiments from Tables 3, 5, 6 and 7. We employ a grid search over the search space. The optimal values are those given in Table 1 and they are used in all of the experiments of Section 3.5.

Remark 3.5. *Both the SELU and ELU activation functions were tested in addition to ReLU but neither performed better than ReLU. We experienced no issues during training with the ReLU activation, perhaps because we are working with quite shallow networks (3 layers) and we do not train for a long time (< 100 iterations). We have intentionally made the software straightforward to adapt such that if, for example, problems with ReLU become apparent in different*

examples, then the activation function can be easily changed.

Stopping rule for training

Rather than train for a fixed number of epochs for each problem, we propose a stopping rule which will determine when to stop training. On a heuristic level, we want to terminate the training procedure if the cost of training for a further epoch outweighs the variance reduction caused by this further training.

Denote the cost of training for one epoch (we assume that this is consistent across epochs) by $\mathcal{C}_{\text{train}}^{(1)}$. Let $\text{Var}\bar{\Gamma}_{\theta_i}$ denote the reduced variance after the i -th epoch with $\text{Var}\bar{\Gamma}_{\theta_0} := \text{Var}\Gamma$. Then, before the i -th epoch, we should decide to stop the training algorithm when

$$\mathcal{C}_{\text{train}}^{(1)} + \Phi^{-1}(1 - \alpha/2)^2 \mathcal{C}_{\text{batch}} \frac{\text{Var}\bar{\Gamma}_{\theta_i}}{\varepsilon^2 S_{\text{batch}}} > \Phi^{-1}(1 - \alpha/2)^2 \mathcal{C}_{\text{batch}} \frac{\text{Var}\bar{\Gamma}_{\theta_{i-1}}}{\varepsilon^2 S_{\text{batch}}}, \quad (3.125)$$

which is equivalent to

$$\text{Var}\bar{\Gamma}_{\theta_{i-1}} - \text{Var}\bar{\Gamma}_{\theta_i} < \frac{\mathcal{C}_{\text{train}}^{(1)}}{K}, \quad (3.126)$$

where

$$K = \frac{\Phi^{-1}(1 - \alpha/2)^2 \mathcal{C}_{\text{batch}}}{\varepsilon^2 S_{\text{batch}}}. \quad (3.127)$$

Of course, the variance after the i -th epoch is unknown until after the epoch, so we estimate the change $\text{Var}\bar{\Gamma}_{\theta_{i-1}} - \text{Var}\bar{\Gamma}_{\theta_i}$ with the change from the previous epoch, that is $\text{Var}\bar{\Gamma}_{\theta_{i-2}} - \text{Var}\bar{\Gamma}_{\theta_{i-1}}$.

In practice, the cost to train for one epoch, $\mathcal{C}_{\text{train}}^{(1)}$, is estimated during training as the running average of the time taken for one epoch of training. The cost to sample one batch, $\mathcal{C}_{\text{batch}}$, is computed before training takes place.

3.5.2 Diffusion models

In the first two experiments (Sections 3.1 - 3.2), we use the explicit Euler scheme (see e.g. [101]), which for a system

$$dX(t) = b(t, X)dt + \sigma(t, X)dW(t), \quad X(s) = x, \quad (3.128)$$

is defined by

$$X_0 = X(s) = x, \quad (3.129)$$

$$X_{k+1} = X_k + b_k h + \sigma_k \Delta_k W, \quad (3.130)$$

where $h > 0$ is a fixed step-size and $\Delta_k W$ are independent Gaussian random variables (or, in the case of many Wiener processes, vectors consisting of independent Gaussian random variables) with zero mean and variance h . The step-size used for each experiment is $h = \frac{T-s}{1000}$. In the experiment of Section 3.3 we use a custom numerical scheme from [99].

Example 3.1 (Geometric Brownian motion). Consider the one-dimensional Black-Scholes model, in which the price $u(t, x)$ of a contingent claim satisfies the Black-Scholes equation

$$\frac{\partial u}{\partial t} + \frac{\sigma^2 x^2}{2} \frac{\partial^2 u}{\partial x^2} + rx \frac{\partial u}{\partial x} - ru = 0 \quad (3.131)$$

with terminal condition

$$u(T, x) = f(x), \quad (3.132)$$

where the short rate $r \in \mathbb{R}$ and the volatility $\sigma > 0$. We consider the case of a call option when $f(x) = (x - K)_+$ for various strikes $K > 0$. The problem (3.131)-(3.132) has the following probabilistic representation:

$$u(s, x) = \mathbb{E} \left[f(X_{s,x}(T)) e^{-r(T-s)} + Z_{s,x}(T) \right], \quad (3.133)$$

where

$$dX(t) = rX(t)dt + \sigma X(t)dW(t), \quad X(s) = x, \quad (3.134)$$

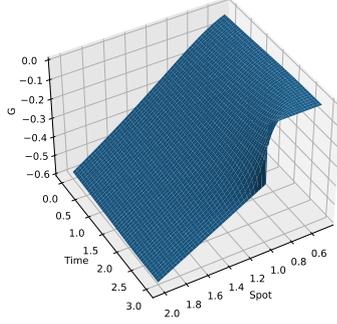
$$dZ(t) = G(t, X(t))e^{-r(t-s)}dW(t), \quad Z(s) = 0. \quad (3.135)$$

Numerical results for various strikes, K , are given in Table 3 for both vanilla MC (i.e. without any control variates) and for our neural variance reduction algorithm. We observe substantial (up to 18 times) speed up of option valuation when the variance reduction is used in comparison with the vanilla MC. The MC tolerance level is set at 10^{-4} in this and all the other experiments, aside from those in Section 3.5.4 where it is set to 10^{-3} . The relative error given in the table (and in all the other tables of Section 3.5.2) corresponds to $\text{Err}_{G_{\theta^*}}$ as defined in (3.85).

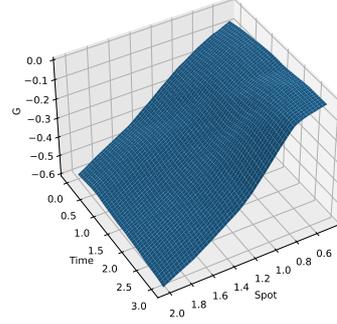
Figure 1 shows the true optimal control variate, as given by Theorem 3.2, as well as the learned approximate control variate in the case that $K = 1$. It can be seen that the learned control variate is similar to the optimal control variate but not especially accurate, indicating that high accuracy is not required for the control variate method to be effective.

Remark 3.6. *If the optimal control variate G^* is known (or approximated well) then this would give knowledge of ∇u through Theorem 3.2, which can be used for computing deltas (e.g., analogously how conditional probabilistic representations accompanied by linear regression is used for computing sensitivities in [11]). However, in this paper we refrain from doing so since we have no guarantee of accuracy of the neural network approximation of G^* , so there is no theoretical guarantee that the learned approximation of ∇u is accurate, and importantly no reliable way to quantify the error. This is in contrast to using G_{θ} as a control variate, since in this case the bias of the neural network approximation is nullified by integrating against the Brownian motion. Moreover, the MC error can be quantified in the normal way by computing the sample variance.*

In the pursuit of computing Greeks, Malliavin Greeks can be considered [44]. Indeed, this was considered in the variance reduction setting in [100]. This may prove to be a fruitful avenue of research but we consider it beyond the scope



(a) Optimal control variate



(b) Learned control variate

Figure 1: The true optimal control variate and the learned control variate for a European call option with strike $K = 1$ under the Black-Scholes model with parameters in Table 3.

of this work.

Table 3: European Call option, $f(x) = (x - K)_+$, under GBM with $r = 0.02$, $\sigma = 0.3$ and $T = 3$: MC approximations (and a 95% confidence interval given after \pm) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	0.39031	0.39043 ± 0.00010	87.2	$1.04 \cdot 10^8$	0.39032 ± 0.00010	5.73	$7.00 \cdot 10^4$	0.034
0.8	0.32826	0.32833 ± 0.00010	80.3	$9.62 \cdot 10^7$	0.32821 ± 0.00010	4.27	$7.00 \cdot 10^4$	0.040
0.9	0.27484	0.27486 ± 0.00010	70.9	$8.49 \cdot 10^7$	0.27478 ± 0.00010	4.31	$7.50 \cdot 10^4$	0.051
1	0.22943	0.22943 ± 0.00010	63.1	$7.56 \cdot 10^7$	0.22940 ± 0.00010	6.17	$1.15 \cdot 10^5$	0.075
1.1	0.19117	0.19115 ± 0.00011	48.4	$5.79 \cdot 10^7$	0.19115 ± 0.00010	6.49	$1.40 \cdot 10^5$	0.098
1.2	0.15914	0.15914 ± 0.00010	46.2	$5.53 \cdot 10^7$	0.15915 ± 0.00010	4.54	$9.00 \cdot 10^4$	0.096
1.3	0.13245	0.13249 ± 0.00010	44.2	$5.29 \cdot 10^7$	0.13247 ± 0.00010	6.48	$1.40 \cdot 10^5$	0.140

Example 3.2 (Multi-dimensional geometric Brownian motion). Consider now the multi-dimensional Black-Scholes model for dynamics of assets' prices. The price $u(t, x)$ of a contingent claim satisfies the PDE:

$$\frac{\partial u}{\partial t} + \frac{\sigma^2}{2} \sum_{i,j=1}^d \rho^{i,j} x^i x^j \frac{\partial^2 u}{\partial x^i \partial x^j} + r \sum_{i=1}^d x^i \frac{\partial u}{\partial x^i} - ru = 0 \quad (3.136)$$

with terminal condition

$$u(T, x) = f(x), \quad (3.137)$$

where $r \in \mathbb{R}$, $\sigma > 0$ and $(\rho)_{ij} = \rho^{i,j}$ is the correlation matrix of the Brownian motion in the following probabilistic representation:

$$u(s, x) = \mathbb{E} \left[f(X_{s,x}(T)) e^{-r(T-s)} + Z_{s,x}(T) \right], \quad (3.138)$$

and

$$dX^i(t) = rX^i(t)dt + \sigma X^i(t)dW_i(t), \quad X^i(s) = x, \quad i = 1, \dots, d, \quad (3.139)$$

$$dZ(t) = G^\top(t, X(t))e^{-r(t-s)}dW(t), \quad Z(s) = 0, \quad (3.140)$$

$W(t) = (W_1(t), \dots, W_d(t))^\top$, and $W_i(t), i = 1, \dots, d$, are correlated Wiener processes. We consider the case of a call-on-max option, that is an option with the payoff

$$f(x) = (\max(x_1, \dots, x_d) - K)_+. \quad (3.141)$$

Table 4 displays the results for Algorithm 1 in the case $d = 3$ with the parameters $r = 0.02, \sigma = 0.3, T = 3, x = 1$ and the correlation coefficients of the Wiener processes:

$$\rho^{1,2} = 0.7, \quad \rho^{1,3} = 0.2, \quad \rho^{2,3} = -0.3.$$

We again see the benefit of using the control variate method which gives speed up of up to 10 times.

Table 4: Call-on-max rainbow option under three-dimensional GBM with $r = 0.02, \sigma = 0.3, T = 3, x = 1$: Monte Carlo approximations (and a 95% confidence interval) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	-	0.73571 ± 0.00010	187	$1.51 \cdot 10^8$	0.73570 ± 0.00010	20.2	$9.15 \cdot 10^5$	0.066
0.8	-	0.64800 ± 0.00010	188	$1.52 \cdot 10^8$	0.64811 ± 0.00010	20.0	$9.20 \cdot 10^5$	0.076
0.9	-	0.56563 ± 0.00010	185	$1.50 \cdot 10^8$	0.56559 ± 0.00010	20.8	$9.65 \cdot 10^5$	0.088
1	-	0.48988 ± 0.00010	177	$1.44 \cdot 10^8$	0.48984 ± 0.00010	17.5	$7.60 \cdot 10^5$	0.091
1.1	-	0.42146 ± 0.00010	155	$1.26 \cdot 10^8$	0.42150 ± 0.00010	20.2	$9.20 \cdot 10^5$	0.120
1.2	-	0.36085 ± 0.00010	153	$1.24 \cdot 10^8$	0.36098 ± 0.00010	18.8	$8.30 \cdot 10^5$	0.130
1.3	-	0.30781 ± 0.00010	134	$1.08 \cdot 10^8$	0.30784 ± 0.00010	17.1	$7.35 \cdot 10^5$	0.140

Example 3.3 (Heston model). Consider the Heston stochastic volatility model [64], under which the price $u(t, x, v)$ of a contingent claim satisfies

the PDE:

$$\frac{\partial u}{\partial t} + \frac{1}{2}x^2v\frac{\partial^2 u}{\partial x^2} + \frac{1}{2}\sigma^2v\frac{\partial^2 u}{\partial v^2} + \sigma\rho vx\frac{\partial^2 u}{\partial x\partial v} + rx\frac{\partial u}{\partial x} + \kappa(\theta - v)\frac{\partial u}{\partial v} - ru = 0, \quad (3.142)$$

$$u(T, x, v) = f(x), \quad (3.143)$$

where $\kappa > 0$, $\theta > 0$, $\sigma > 0$, $\rho \in (-1, 1)$ and $r \in \mathbb{R}$, such that $2\kappa\theta > \sigma^2$. The terminal condition, $f(x)$, corresponds to the payoff function. We consider the case of a European call option where

$$f(x) = (x - K)_+, \quad (3.144)$$

for some strike price $K > 0$. The Heston PDE problem (3.142)-(3.143) has a probabilistic representation of the form:

$$u(s, x, v) = \mathbb{E}\left[f(X_{s,x,v}(T))e^{-r(T-s)} + Z_{s,x,v}(T)\right], \quad (3.145)$$

$$dX(t) = rX(t)dt + \sqrt{V(t)}X(t)dW_1(t), \quad (3.146)$$

$$dV(t) = \kappa(\theta - V(t)) + \sigma\sqrt{V(t)}\left(\rho dW_1(t) + \sqrt{1 - \rho^2}dW_2(t)\right), \quad (3.147)$$

$$dZ(t) = G(t, X(t), V(t))e^{-r(t-s)}dW(t), \quad (3.148)$$

$X(s) = x$, $V(s) = v$, $Z(s) = 0$ and where $W(t) = (W_1(t), W_2(t))^T$ is a two-dimensional standard Wiener process.

To simulate trajectories of (3.146)-(3.147), we use the explicit Euler scheme for X and the fully implicit Euler scheme for V [99]. For a step-size $h > 0$,

$$X_{k+1} = X_k + rX_k h + \sqrt{V_k}X_k\Delta_k W, \quad (3.149)$$

$$V_{k+1} = V_k + \kappa(\theta - V_{k+1})h - \frac{\sigma^2}{2}h + \sigma\sqrt{V_{k+1}}\Delta_k W. \quad (3.150)$$

The importance of using the fully implicit Euler scheme for V lies in the fact that the explicit Euler scheme does not necessarily preserve positivity of V , while the semi-implicit scheme (3.149)-(3.150) guarantees positivity

of V under $\kappa\theta \geq \sigma^2/2$ [99]. Note that the implicitness in (3.150) can be resolved analytically by solving the quadratic equation. The results are presented in Table 5, which demonstrates efficiency of the proposed control variate method, it is up to 42 times faster than vanilla MC.

Table 5: European Call option under the Heston model with $v = 0.15$, $r = 0.02$, $\kappa = 0.25$, $\theta = 0.5$, $\sigma = 0.3$, $\rho = -0.3$, $T = 3$: MC approximations (and a 95% confidence interval) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	0.47517	0.47520 ± 0.00010	1360	$3.09 \cdot 10^8$	0.47519 ± 0.00008	45.3	$2.58 \cdot 10^6$	0.14
0.8	0.42623	0.42625 ± 0.00011	1120	$2.57 \cdot 10^8$	0.42636 ± 0.00010	30.9	$1.65 \cdot 10^6$	0.16
0.9	0.38271	0.38274 ± 0.00010	1180	$2.72 \cdot 10^8$	0.38267 ± 0.00011	32.0	$1.73 \cdot 10^6$	0.19
1	0.34406	0.34401 ± 0.00010	1070	$2.47 \cdot 10^8$	0.34407 ± 0.00009	54.6	$3.17 \cdot 10^6$	0.24
1.1	0.30977	0.30977 ± 0.00010	1170	$2.7 \cdot 10^8$	0.30971 ± 0.00010	27.4	$1.43 \cdot 10^6$	0.20
1.2	0.27934	0.27927 ± 0.00009	1250	$2.88 \cdot 10^8$	0.27927 ± 0.00011	39.5	$2.20 \cdot 10^6$	0.28
1.3	0.25232	0.25230 ± 0.00010	990	$2.28 \cdot 10^8$	0.25232 ± 0.00010	46.3	$2.61 \cdot 10^6$	0.34

3.5.3 Models with a non-singular Lévy measure

In this section and the subsequent section (Section 3.5.4), we use the restricted jump-adapted numerical integration scheme from [36, Algorithm 1]. Here we give its brief description for completeness.

For the SDEs

$$dX = b(t, X) - F(t, X)\gamma_\epsilon dt + \sigma(t, X)dw(t) \quad (3.151)$$

$$+ F(t, X)\beta_\epsilon dW(t) + \int_{|z| \geq \epsilon} F(t, X)zN(dt, dz), \quad (3.152)$$

$$X(s) = x, \quad (3.153)$$

we set $X_0 = x$ and obtain the approximation X_{k+1} from X_k as follows. We find the next time-step $\theta = \delta \wedge h$, where $h > 0$ is the pre-defined maximum step-size and δ is the time to the next jump sampled with the intensity $\lambda_\epsilon = \int_{|z| \geq \epsilon} \nu(z)$. Then, if $\theta = h$, we use the standard explicit Euler scheme with no jumps. If $\theta < h$, we use

$$X_{k+1} = X_k + (b_k - F_k)\gamma_\epsilon\theta + \sigma_k\Delta_k w + F_k\beta_\epsilon\Delta W_k + F_k J, \quad (3.154)$$

where $\Delta_k w$ and $\Delta_k W$ are the respective Brownian increments over the

step θ and J is the size of the jump, sampled (independently of other jumps, δ , $\Delta_k w$ and $\Delta_k W$) according to the density

$$\rho_\varepsilon(z) := \frac{\nu(z)\mathbb{1}_{|z|>\varepsilon}}{\lambda_\varepsilon}. \quad (3.155)$$

The approximations for Y and Z are simulated using the explicit Euler scheme, with the same random time step θ as X . The (maximum) step-size used in all the experiments is $h = \frac{T-s}{1000}$.

Example 3.4 (Merton model). Consider the one-dimensional Merton jump-diffusion model [94], under which the price $u(t, x)$ of a contingent claim satisfies the PIDE:

$$\begin{aligned} \frac{\partial u}{\partial t}(t, x) + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 u}{\partial x^2}(t, x) + (r - \beta\lambda)x \frac{\partial u}{\partial x}(t, x) - ru(t, x) \\ + \frac{\lambda}{\sqrt{2\pi\gamma^2}} \int_{\mathbb{R}} [u(t, xe^z) - u(t, x)] \exp\left\{-\frac{(z - \alpha^2)^2}{2\sigma^2}\right\} dz = 0, \end{aligned} \quad (3.156)$$

where $r \in \mathbb{R}$, $\sigma > 0$, $\lambda > 0$, $\alpha, \gamma \in \mathbb{R}$ and $\beta = \exp(\alpha + \frac{1}{2}\gamma^2) - 1$. The terminal condition is given by

$$u(T, x) = f(x). \quad (3.157)$$

The probabilistic representation of (3.156)-(3.157) takes the form:

$$u(s, x) = \mathbb{E}\left[f(X_{s,x}(T))e^{-r(T-s)} + Z_{s,x}(T)\right] \quad (3.158)$$

with

$$dX(t) = X(t-)((r - \lambda\beta)dt + \sigma dW(t) + J(t)dN(t)), \quad (3.159)$$

$$\begin{aligned} dZ(t) = G_W(t, X(t))e^{-r(t-s)}dW(t) + G_N(t-, X(t-))e^{-r(t-s)}J(t)dN(t) \\ - \lambda e^{-r(t-s)} \int_{\mathbb{R}} G_N(t, X)z \exp\left\{-\frac{(z - \alpha^2)^2}{2\sigma^2}\right\} dz dt, \end{aligned} \quad (3.160)$$

with $X(s) = x$, $Z(s) = 0$, where $W(t)$ is a one-dimensional standard Wiener process and $N(t)$ is a Poisson process with intensity λ . The jumps

have a shifted log-normal distribution:

$$J_i = \exp(\eta_i) - 1, \quad i = 1, 2, \dots$$

with $\eta_i \sim \mathcal{N}(\alpha, \gamma^2)$. The mean jump size is $\mathbb{E}J_1 =: \beta = \exp(\alpha + \frac{1}{2}\gamma^2) - 1$.

The price of a European call option, $u(t, x)$, has the terminal condition

$$f(x) = (x - K)_+,$$

for a strike price $K > 0$.

Table 6 shows the results of Algorithm 2 in the case of $r = 0.02$, $\sigma = 0.2$, $\lambda = 1$, $\alpha = -0.05$, $\gamma = 0.3$. The proposed control variates method here is up to 30 times faster than vanilla MC. The relative error given in the table (and in all the other tables in the next subsections) corresponds to $\text{Err}_{G_{\theta^*}}$ as defined in (3.123).

Table 6: European Call option under Merton model with $r = 0.02$, $\sigma = 0.2$, $\lambda = 1$, $\alpha = -0.05$, $\gamma = 0.3$ and $T = 3$: MC approximations (and a 95% confidence interval) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	0.41361	0.41346 ± 0.00010	1860	1.61·10 ⁸	0.41364 ± 0.00010	62.1	1.40·10 ⁶	0.15
0.8	0.35593	0.35575 ± 0.00009	2010	1.74·10 ⁸	0.35592 ± 0.00011	85.5	2.05·10 ⁶	0.21
0.9	0.30592	0.30575 ± 0.00011	1480	1.28·10 ⁸	0.30604 ± 0.00010	104	2.49·10 ⁶	0.26
1	0.26298	0.26278 ± 0.00010	1460	1.27·10 ⁸	0.26302 ± 0.00010	127	3.09·10 ⁶	0.34
1.1	0.22634	0.22624 ± 0.00011	1200	1.05·10 ⁸	0.22633 ± 0.00010	140	3.47·10 ⁶	0.42
1.2	0.19519	0.19502 ± 0.00011	1090	9.52·10 ⁷	0.19532 ± 0.00010	151	3.79·10 ⁶	0.51
1.3	0.16877	0.16866 ± 0.00011	949	8.28·10 ⁷	0.16881 ± 0.00010	167	4.17·10 ⁶	0.62

3.5.4 Models with a singular Lévy measure

In this section, we test Algorithm 2 on SDEs driven by Lévy processes with infinite activity of jumps. In this case the algorithm is of particular importance because variance of quantities of interest is typically very large and practical use of such models in financial engineering requires efficient variance reduction.

Consider the process to model an underlier:

$$S_i(t) = S_i(s) \exp \{rt + X_i(t)\}, \quad (3.161)$$

where $r > 0$ and X is a d -dimensional process defined as

$$\begin{aligned} X(T) = & \int_s^T b(t, X(t-)) dt + \int_s^T \sigma(t, X(t-)) dw(t) \\ & + \int_s^T \int_{\mathbb{R}} F(t, X(t-)) z \hat{N}(dt, dz), \end{aligned} \quad (3.162)$$

with Lévy measure given by

$$\nu(dz) = \begin{cases} C_- e^{-\mu(|z|-1)} dz & \text{if } z < -1, \\ C_- |z|^{-(\alpha+1)} dz & \text{if } -1 \leq z < 0, \\ C_+ |z|^{-(\alpha+1)} dz & \text{if } 0 < z \leq 1, \\ C_+ e^{-\mu(|z|-1)} dz & \text{if } 1 < z, \end{cases} \quad (3.163)$$

where C_- , C_+ and μ are positive constants and $\alpha \in (0, 2)$. Throughout, we take $\sigma(t, x)$ to be a constant matrix, $\sigma(t, x) = \sigma \in \mathbb{R}^{d \times d}$ and $F(t, x) = (f_1, \dots, f_d) \in \mathbb{R}^d$. Since the discounted price processes, $\tilde{S}_i(t) = e^{-rt} S_i(t)$, should be martingales, the drift component $b(t, x)$ is chosen as (see [36] or [3, Sec 5.2])

$$b_i = -\frac{1}{2} \sum_{j=1}^d \sigma_{ij}^2 - \int_{\mathbb{R}} (e^{f_i z} - 1 - f_i z \mathbb{1}_{|z| < 1}) \nu(dz). \quad (3.164)$$

We consider this model in one and four dimensions. In all of the examples, we choose $C_- = C_+ = 1$, $\alpha = 0.5$, $\mu = 2$ and $\varepsilon = 10^{-3}$.

Example 3.5 (One-dimensional European call). We consider the problem of pricing a European call option under model (3.161)-(3.164) with $d = 1$. Table 7 shows the results for Algorithm 2 in the case of $r = 0.02$, $\sigma_{11} = 0.2$ and $f_1 = 0.2$. The computational speed up achieved here by the proposed control variate method is up to 8 times in comparison with vanilla MC.

Example 3.6 (Four-dimensional call-on-max option). As before, consider the model (3.161)-(3.164) with $d = 4$ and the terminal condition

$$f(x) = (\max(x_1, x_2, x_3, x_4) - K)_+ \quad (3.165)$$

Table 7: European Call option under exponential Lévy model (3.161): MC approximations (and a 95% confidence interval) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	-	0.46285±0.00100	225	3.00·10 ⁶	0.46298±0.00095	28.1	1.30000·10 ⁵	0.38
0.8	-	0.41284±0.00101	208	2.79·10 ⁶	0.41376±0.00103	27.8	1.25000·10 ⁵	0.45
0.9	-	0.36926±0.00099	205	2.76·10 ⁶	0.36934±0.00101	25.2	1.40000·10 ⁵	0.52
1	-	0.33181±0.00100	193	2.58·10 ⁶	0.33189±0.00100	29.4	1.35000·10 ⁵	0.57
1.1	-	0.29821±0.00102	175	2.34·10 ⁶	0.29721±0.00098	25.7	1.25000·10 ⁵	0.59
1.2	-	0.26838±0.00102	166	2.22·10 ⁶	0.26841±0.00101	27.3	1.40000·10 ⁵	0.72
1.3	-	0.24385±0.00105	148	1.99·10 ⁶	0.24272±0.00105	25.4	1.45000·10 ⁵	0.84

for some $K > 0$. Table 8 shows the results of Algorithm 2 in the case of $r = 0.02$, $f = (0.2, 0.15, 0.15, 0.1)^\top$ and

$$\sigma = 0.15L, \quad (3.166)$$

where L is the lower triangular matrix obtained via the Cholesky decomposition of the correlation matrix

$$LL^\top = \begin{bmatrix} 1 & 0.87 & 0.94 & 0.86 \\ 0.87 & 1 & 0.87 & 0.93 \\ 0.94 & 0.87 & 1 & 0.96 \\ 0.86 & 0.93 & 0.96 & 1 \end{bmatrix}. \quad (3.167)$$

Here the achieved speed up is up to 14 times.

Table 8: Rainbow call-on-max option under four-dimensional exponential Lévy model (3.161): Monte Carlo approximations (and a 95% confidence interval) with and without a control variate.

K	$u(0, 1)$	Vanilla MC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	-	0.61422±0.00100	699	2.28·10 ⁶	0.61501±0.00101	47.9	1.00·10 ⁵	0.26
0.8	-	0.53704±0.00099	699	2.28·10 ⁶	0.53739±0.00080	72.9	1.65·10 ⁵	0.31
0.9	-	0.46754±0.00100	650	2.12·10 ⁶	0.46939±0.00103	51.1	1.05·10 ⁵	0.36
1	-	0.40659±0.00100	623	2.03·10 ⁶	0.40621±0.00100	50.3	1.10·10 ⁵	0.42
1.1	-	0.35247±0.00101	574	1.87·10 ⁶	0.35399±0.00098	60.2	1.25·10 ⁵	0.50
1.2	-	0.30811±0.00102	530	1.72·10 ⁶	0.30763±0.00097	56.9	1.25·10 ⁵	0.57
1.3	-	0.26923±0.00099	510	1.68·10 ⁶	0.26925±0.00100	61.6	1.40·10 ⁵	0.71

3.5.5 Comparison with alternative complexity and variance reduction methods

So far, we have only compared our neural control variate method to vanilla MC. In this section, we compare our method to other standard complexity and variance reduction techniques. We discuss the Multi-level Monte Carlo (MLMC) method [55] and also a crude control variate method [57].

Example 3.7 (Multilevel Monte Carlo method). Let us consider the problem of pricing a call-on-max option on the model (3.161)-(3.164) with $d = 2$. That is,

$$f(x) = (\max(x_1, x_2) - K)_+ \quad (3.168)$$

for $K > 0$.

In contrast to the previous experiments, here we do not compare our method with the vanilla MC method but instead the MLMC method. We fix $h = T/1024$, and then choose the levels of the MLMC method to be $\{T/N : N = 16, 64, 256, 1024\}$, i.e. geometrically decreasing step size by a factor of 4. We use the jump-adapted Euler scheme, where between each jump the Brownian component is approximated by the Euler scheme while the jumps are simulated exactly [112]. In the context of MLMC, the jumps happen at the same time and with the same size across the coarse and fine paths [55]. Table 9 details the results for a range of strike prices. We observe a speed-up of up to 5.5 times achieved by our neural control variate method compared to MLMC.

Remark 3.7. *We experimentally observed that MLMC outperforms our neural control variate method in the case of standard diffusions considered in Sections 3.1, 3.2 and 3.3 (usually providing a 1.5-3 times speed-up). However, in the case of jump-diffusions the neural control variate method is competitive and performs better when the jump rate is high. In these cases, MLMC is less effective since the advantage of having a larger time-step is diminished when jumps are frequent. We also expect that the neural control variate method can outperform MLMC when stochastic models are more complex and hence cannot be simulated*

Table 9: Rainbow call-on-max option under two-dimensional exponential Lévy model (3.161): Monte Carlo approximations (and a 95% confidence interval) with a control variate and MLMC approximation.

K	$u(0, 1)$	MLMC			Control Variate MC			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	-	0.71275±0.00049	94.4	-	0.71304±0.00050	17.0	9.50·10 ⁴	0.11
0.8	-	0.65320±0.00049	91.8	-	0.65331±0.00050	20.0	1.25·10 ⁵	0.14
0.9	-	0.59884±0.00049	89.2	-	0.59901±0.00049	21.9	1.30·10 ⁵	0.15
1	-	0.54941±0.00049	86.3	-	0.54955±0.00050	21.8	1.25·10 ⁵	0.17
1.1	-	0.50511±0.00049	82.8	-	0.50456±0.00049	24.0	1.65·10 ⁵	0.20
1.2	-	0.46439±0.00049	80.0	-	0.46466±0.00049	24.3	1.75·10 ⁵	0.23
1.3	-	0.42781±0.00049	77.4	-	0.42757±0.00050	28.9	2.05·10 ⁵	0.27

with larger time steps (due to some stability restrictions) as required for MLMC efficiency.

Example 3.8 (Crude control variate). We now compare our method to using the well-known method of employing the terminal spot value as a control variate [57]. That is, for an asset price process $X(t)$, using the random variable

$$\Gamma = e^{-rT} f(X(T)) + c(e^{-rT} X(T) - X(0)), \quad (3.169)$$

for some $c \in \mathbb{R}$. This has the same expectation as $e^{-rT} f(X(T))$ given that the discounted process $e^{-rt} X(t)$ is a martingale under the pricing measure.

We work in the same setting as Section 3.3, pricing a European call under the Heston model (3.142)-(3.143). The parameter $c \in \mathbb{R}$ is chosen in the normal way to minimize the variance of Γ , see e.g. [57]. Table 10 compares the results of this method to Algorithm 1. We see a speedup across all strikes, but particularly for out-of-money options. It is well known that the terminal spot control variate performs worse for these options, see [57].

3.5.6 Transfer learning

In the previous experiments, we initialize the weights of the ANN each time the parameters of the financial model change. However, if the change in parameters is small, e.g. when we vary the strike price while keeping

Table 10: European call option under the Heston model with $v = 0.15$, $r = 0.02$, $\kappa = 0.25$, $\theta = 0.5$, $\sigma = 0.3$, $\rho = -0.3$, $T = 3$: Monte Carlo approximations (and a 95% confidence interval) with our neural control variate and a crude control variate.

K	$u(0, 1)$	Crude Control Variate			Neural Control Variate			Relative Error
		$\hat{u}(0, 1)$	Time (s)	M	$\hat{u}(0, 1)$	Time (s)	M	
0.7	0.47517	0.47524 ± 0.00010	57.6	$9.60 \cdot 10^6$	0.47519 ± 0.00008	45.3	$2.58 \cdot 10^6$	0.14
0.8	0.42623	0.42633 ± 0.00010	77.2	$12.9 \cdot 10^7$	0.42636 ± 0.00010	30.9	$1.65 \cdot 10^6$	0.16
0.9	0.38271	0.38279 ± 0.00010	97.3	$16.4 \cdot 10^7$	0.38267 ± 0.00011	32.0	$1.73 \cdot 10^6$	0.19
1	0.34406	0.34415 ± 0.00010	118	$19.9 \cdot 10^7$	0.34407 ± 0.00009	54.6	$3.17 \cdot 10^6$	0.24
1.1	0.30977	0.30983 ± 0.00010	138	$23.2 \cdot 10^7$	0.30971 ± 0.00010	27.4	$1.43 \cdot 10^6$	0.20
1.2	0.27934	0.27934 ± 0.00010	157	$26.3 \cdot 10^7$	0.27927 ± 0.00011	39.5	$2.20 \cdot 10^6$	0.28
1.3	0.25232	0.25229 ± 0.00010	174	$29.2 \cdot 10^7$	0.25232 ± 0.00010	46.3	$2.61 \cdot 10^6$	0.34

the other parameters fixed, we can use the previous weights of the ANN as the initial weights of the next ANN. This approach is termed as transfer learning (see e.g. [59]). Such a procedure can reduce computational costs further by decreasing the training time. We demonstrate this with the following example. Consider the same experiment as in Section 3.2 (see Table 4). Table 11 shows the time taken for the control variate method from Table 4, where the weights of the ANN are initialized each time, compared to the method of transfer learning where the weights are only initialized once. Note that there is no difference in time for $K = 0.7$, since the weights are initialized in the same way. We see that using transferred weights can give up to 2 times of further speed up, giving overall speed-up up to 20 times in comparison with the plain vanilla MC. Hence, transfer learning can further accelerate variance reduction offered by Algorithms 1 and 2.

3.6 Conclusion and future research

In this chapter, we considered variance reduction methods for Monte Carlo simulation of a SDE functionals. We extended the theoretical results regarding control variates and importance sampling to the case that the SDE is driven by Lévy noise. We allowed for the case of infinite activity by approximating small jumps with a Brownian motion with appropriate covariance structure. We use this theoretical result to motivate a two-pass algorithm in which we train a neural network to minimize the empirical

Table 11: Experiment from Section 3.2: comparison of times using weights re-initialized and weights transferred from previous simulations with different K .

K	Time using new weights	Time using transferred weights
0.7	20.2	20.2
0.8	20.0	11.5
0.9	20.8	10.2
1	17.5	8.6
1.1	20.2	8.3
1.2	18.8	11.0
1.3	17.1	8.1

variance of the functional before sampling while using the neural network as a control variate. The fact that the neural network is integrated against the driving noise means that no bias is introduced via the approximation error of the neural network. We demonstrated the effectiveness of the algorithms on a range of problems from option pricing, which is a natural application of the Monte Carlo evaluation of SDE functionals.

We mention two possible areas of future research. The first is the consideration of Dirichlet problem and the extension of theoretical results to this case. This is considered, for Brownian motion only, in e.g. [101]. Financially, this corresponds to the pricing of barrier options, so there is a clear application for this research. The second area of research is towards the combination of MLMC and the algorithms proposed in this chapter. For instance, the first-pass (the training of the control variates) can be done in the usual manner, while the second pass can utilize the MLMC method. Such research could result in further efficiency but was beyond the scope of this chapter.

4 Well-posedness and approximation of reflected McKean-Vlasov SDEs

In this chapter, we study reflected McKean-Vlasov stochastic differential equations (SDEs) and their particle approximations in smooth bounded domains. Well-posedness is established in a general setting under some mild assumptions on the coefficients for both the particle system and its mean-field limit. We also consider reflected consensus-based optimization algorithms, which can be formulated in continuous-time by reflected McKean-Vlasov SDEs. We prove convergence of the particle system to the mean-field limit with optimal rate. For the consensus-based optimization, we prove convergence of the mean-field limit to the global minimum of the optimization problem. Several numerical examples from constrained optimization are presented. This chapter is based on the paper [65]:

- Hinds, P. D., Sharma, A., and Tretyakov, M. V. “Well-posedness and approximation of reflected McKean-Vlasov SDEs with applications”. *Mathematical Models and Methods in Applied Sciences* 35(08) (2025), pp. 1845–1887.

but will have additional exposition.

4.1 Introduction

Reflected stochastic differential equations (SDEs) are used to model processes confined to a domain with boundary, where the solution is reflected in a prescribed direction when it hits the boundary. At the same time, McKean-Vlasov SDEs have coefficients with non-linear dependencies on the law of the solution. This chapter is devoted to the study of reflected McKean-Vlasov SDEs which can model systems with constraints and mean-field interactions. We illustrate the practical relevance of reflected mean-field SDEs via the problem of constrained global optimization.

Optimization problems are ubiquitous in the applied sciences and have found resurgence due to advancements in machine learning. While

convex optimization is well-studied [107], the task of non-convex global optimization poses additional challenges [69]. *Metaheuristic methods* are a popular class of methods which have been used to numerically solve global optimization problems. Such methods consist of a high-level algorithmic framework to coordinate low-level heuristics in order to efficiently explore and exploit the solution space. Examples include particle swarm optimization [75] and differential evolution [128] among others. While there is typically limited theoretical foundation for such models, these methods have been found to be effective in practice [141].

Interacting particle system based models have resulted in better understanding of many natural and social phenomena [62, 10, 9, 29]. Therefore, it is not surprising that interacting particle based methods for optimization and sampling have also gained traction (see [87, 28, 91, 49, 50, 115, 71]) due to their enhanced capability to explore a non-convex energy landscape. These systems are driven by SDEs and hence the tools from stochastic calculus, in particular mean-field theory [130] and stochastic numerics [101] are available to establish their convergence.

Let us discuss the contributions of this chapter along with the comparison with existing literature. In the chapter, we first establish existence and uniqueness of reflected McKean-Vlasov SDEs and their particle approximations in a bounded and non-convex domain setting (Section 4.4). The well-posedness was shown in a convex domain setting in [1] and in a non-convex and bounded setting, but allowing only first order interaction, in [131]. In [144] well-posedness was also considered in a non-convex domain setting but for McKean-Vlasov SDEs with singular drift and extra conditions on the diffusion coefficient.

Our next result is convergence of the interacting particle system towards its mean-field limit with the optimal rate of convergence (Section 4.5). We provide two example models for consensus-based optimization, a metaheuristic method introduced in [111], for which we establish convergence to the global minimum (Section 4.6).

Consensus-based optimization (CBO) is a metaheuristic method

which was introduced in [111]. In these methods, the energy landscape of the objective function is explored by N interacting particles. Each particle's position updates iteratively by drifting towards a consensus point, a weighted average of particle locations according to how well they minimize the objective function. The particles also explore the space through random perturbations. These dynamics allow for the continuous-time formulation of the CBO model via SDEs. The mathematical framework for CBO has been developed in [28, 30, 61, 71].

One recent direction of research relating to CBO has been towards constrained CBO; that is, using CBO models for constrained optimization problems. We briefly mention other strategies which have been proposed and employed to handle constraints in the literature. In [18], a CBO model is considered where a penalty function is added to the objective function f , which penalizes the objective when points are outside the feasible set \bar{G} (f is defined on \mathbb{R}^d). An Euler scheme is applied to simulate the corresponding particle system. Convergence of the algorithm to the optimal solution is shown, both when the exact penalty parameter ϵ is known, and also when it is iteratively updated. In [31], like [18], the authors also use a penalized objective function presenting several numerical experiments. In [6], a projection technique is considered for discrete time CBO model driven by common noise, but no continuous-time analogue is considered.

The distinction of our work is that we encode the constraints directly into the stochastic dynamics with the use of reflected SDEs. Hence, we do not need to modify the objective function, noting that modifying f so that it preserves the global optimum within/near the constraint set \bar{G} is not a trivial task, especially in high dimensions. Moreover, introducing artificial barriers in the objective function typically leads to SDEs with their coefficients taking very large values outside the domain \bar{G} (effectively, making the SDEs very stiff) which in turn requires use of numerical methods approximating these stiff SDEs with small time steps or of complex nature. This problem does not arise within the reflected

SDEs setting considered in this work.

Further, we can split our analysis into two steps. In the first step, we establish convergence of continuous-time dynamics to the global optimum. The next step is to approximate these continuous-time reflected dynamics, with a discrete-time system; for which a plethora of numerical discretization schemes are available to us, see e.g. [101]. Each of these schemes can be considered as an optimization technique. For mean-square (strong) approximation, which is required in this setting, we have the projection scheme [109, 126] and the penalty scheme [126].

This chapter is structured as follows. In Section 4.2, we provide preliminary material regarding reflected SDEs, McKean-Vlasov SDEs and CBO. In Section 4.3, we introduce the set-up of reflected McKean-Vlasov SDEs and describe and introduce two models for CBO. We establish well-posedness in Section 4.4 and prove propagation of chaos in Section 4.5. Convergence of the CBO models to the global minimum of the optimization problem is shown in Section 4.6. Finally, Section 4.7 contains several numerical experiments to test the performance of the CBO models on constrained optimization problems.

4.2 Preliminaries

In this section, we provide background material specific to this chapter, which builds on the preliminary material established in Chapter 2. We introduce both reflected SDEs and McKean-Vlasov SDEs, the combination of which is the object of study in later sections of this chapter. We also discuss consensus-based optimization methods.

4.2.1 Reflected stochastic differential equations

In many real-world systems, the state of a process is constrained to lie within a specific domain. For example, in physical systems particles may be confined to a container [116]; or in optimization problems solutions must lie within a feasible region [22]. Standard SDEs do not account for such constraints, which can lead to solutions that violate the

domain boundaries. To address this, reflected SDEs were introduced. These equations incorporate a reflection mechanism that ensures the solution remains within the domain by forcing it back whenever it hits the boundary.

Reflected processes were first studied by Skorokhod [125], focusing on one-dimensional diffusions restricted to $[0, \infty)$. The term *Skorokhod problem* is now used for the associated deterministic problem of finding, for a given path, a correction that keeps it in a domain. In higher dimensions, the works of Tanaka [132] and Lions and Sznitman [90] established well-posedness results for reflected SDEs in both convex domains and smoothly bounded domains. Further work has been carried out in, e.g., [119, 40, 41].

Throughout, we will use $G \subset \mathbb{R}^d$ to denote a bounded domain and ∂G to denote its boundary, while $\bar{G} = G \cup \partial G$ will denote its closure.

The Skorokhod Problem

Let us begin in the deterministic setting where we are given an unconstrained (continuous) path $\psi : [0, T] \rightarrow \mathbb{R}$, with the simple case of $G = (0, \infty)$. We would like to find the pair of continuous functions (ϕ, η) which satisfy:

1. $\phi(t) = \psi(t) + \eta(t)$ for $t \in [0, T]$,
2. $\phi(t) \geq 0$ for $t \in [0, T]$,
3. η is non-decreasing with $\eta(0) = 0$,
4. $\int_0^t \mathbb{1}_{\{\phi(s)=0\}} d\eta(s) = \eta(t)$ for $t \in [0, T]$.

The above conditions are quite natural. The process η acts as an accumulation of forces which drive the unconstrained process ψ into a constrained process ϕ . The final condition ensures that this forcing action only occurs when ϕ is on the boundary $\partial G = \{0\}$.

This problem has a unique solution given by

$$\eta(t) = - \min_{s \in [0, t]} (\psi(s) \wedge 0), \quad (4.1)$$

$$\phi(t) = \psi(t) + \eta(t), \quad (4.2)$$

see [125, 110].

There are two ways to generalize from here. First, instead of the half-space, we can consider a general domain $G \subset \mathbb{R}^d$. Second, we can leave the deterministic setting and consider the case that we have dynamics driven by an SDE. For the treatment of both cases separately, one can consult e.g. Pilipenko [110]. Here, we jump straight to the general case of an SDE constrained to a domain $G \subset \mathbb{R}^d$.

We will consider the case of a bounded domain $G \subset \mathbb{R}^d$, with smooth boundary ∂G . The exact smoothness condition will be specified later, but for now assume it is smooth enough to allow for a unique-valued inward normal direction, ν , on ∂G .

Let $W(t)$, $t \geq 0$, be a standard d -dimensional Wiener process defined on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$. For measurable functions $b : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : [0, \infty) \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, we say that the pair of continuous and \mathcal{F}_t -adapted processes (X, L) is the solution to the reflected SDE

$$dX(t) = b(t, X(t))dt + \sigma(t, X(t))dW + \nu(X(t))dL(t), \quad (4.3)$$

with initial condition $X(0) = X_0$ if

1. $X(t) \in \bar{G}$ for all $t \geq 0$;
2. L is non-decreasing with $L(0) = 0$ and for all $t \geq 0$,

$$\int_0^t \mathbb{1}_{\{X(s) \in \partial G\}} dL(s) = L(t); \quad (4.4)$$

3. For all $t \geq 0$,

$$\begin{aligned} X(t) = X_0 &+ \int_0^t b(s, X(s)) \mathbf{d}s + \int_0^t \sigma(s, X(s)) \mathbf{d}W(s) \\ &+ \int_0^t \nu(X(s)) \mathbf{d}L(s), \end{aligned} \quad (4.5)$$

provided all the integrals are well-defined. Under a Lipschitz assumption on the coefficients b and σ , existence and uniqueness of a solution (X, L) is proved in [90] for smooth bounded domains, and in [132] for convex domains.

4.2.2 McKean-Vlasov stochastic differential equations

A McKean-Vlasov SDE is a generalization of an SDE where the coefficients are allowed to depend on the current distribution of the solution, as well as on the current state of the process. They are of the form

$$\mathbf{d}X(t) = b(t, X(t), \mathcal{L}_{X(t)}) \mathbf{d}t + \sigma(t, X(t), \mathcal{L}_{X(t)}) \mathbf{d}W(t), \quad (4.6)$$

with a given initial condition $X(0) = X_0$. Here, $\mathcal{L}_{X(t)}$ is the distribution of $X(t)$. McKean-Vlasov SDEs arose in the field of statistical physics [70, 93], and have since been adopted as a framework in probability, in which they exemplify non-linear Markov processes. They have found application in mean-field games [86], systematic risk modelling [27], and control theory [25], as well as in other areas.

Let $W^i(t)$, $t \geq 0$, $i = 1, \dots, N$ be independent Wiener processes on $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$. McKean-Vlasov SDEs arise as the limit of interacting particle systems of the form

$$\mathbf{d}X^i(t) = b(t, X^i(t), \hat{\mu}_{X(t)}) \mathbf{d}t + \sigma(t, X^i(t), \hat{\mu}_{X(t)}) \mathbf{d}W^i(t), \quad (4.7)$$

$i = 1, \dots, N$, where $\hat{\mu}_X$ denotes the empirical measure of the particles:

$$\hat{\mu}_{X(t)} := \frac{1}{N} \sum_{i=1}^N \delta_{X^i(t)}. \quad (4.8)$$

Here, δ_x denotes the Dirac delta at x . Particle systems of this form naturally arise from systems with pairwise interactions between particles. As the number of individual particles, N , tends to infinity, it becomes unnecessary (and hopeless, in any case) to keep track of the location, dynamics and interactions of each individual particle and one can instead study a single particle and its interactions with its distribution $\mathcal{L}_{X(t)}$.

The main theoretical topics of interest relating to McKean-Vlasov SDEs are the well-posedness of them and their interacting particle systems, the propagation of chaos, and numerical simulation of the particle system. These topics have been studied in [130, 102, 33, 85, 26, 20, 117] and references therein.

4.2.3 Consensus-based optimization

Starting in [111], the metaheuristic method of consensus-based optimization (CBO) was considered. In this model, the energy landscape of the objective function is explored by N particles. Each particle broadcasts its current location to the other particles via an average location which is weighted according to the energy landscape: particles where the objective function is small are given higher weights than those where the objective function is large. The position of each particle updates iteratively such that each particle drifts towards this weighted average, whilst also exploring its current neighbourhood through random perturbations. These dynamics allow for the continuous-time formulation of the CBO model via stochastic differential equations.

Each particle's location X^i , $i = 1, \dots, N$, changes throughout time and adheres to the SDE:

$$dX^i(t) = -\beta(X^i(t) - \bar{X}(t))dt + \sigma \text{Diag}(X^i(t) - \bar{X}(t))dW^i(t), \quad (4.9)$$

with $X^i(0) = X_0^i$ and where $\beta > 0$ and $\sigma > 0$ are hyperparameters which control the drift and diffusion strength. The value $\bar{X}(t)$ is called the

particles' consensus at time t and defined by

$$\bar{X} := \bar{X}^{f,\alpha} := \frac{\sum_{i=1}^N X^i e^{-\alpha f(X^i)}}{\sum_{i=1}^N e^{-\alpha f(X^i)}}, \quad (4.10)$$

for some parameter $\alpha > 0$. For large α , the consensus \bar{X} acts as a proxy for the particle which performs best on the objective function, i.e. $\arg \min_{i=1,\dots,N} f(X^i)$. In machine learning, the weighting function is called soft-max and is a smooth way of approximating the argmax function (or in this case argmin). Indeed, the reason we use the particles' consensus rather than the current best particle is so that we can perform analysis on the SDE and its mean-field limit.

The dynamics of the particles, Equation (4.9), has two parts. The first part, due to the drift term, causes each particle to drift towards the consensus. This is exploiting behaviour: each individual particle should want to do better than where it currently is, and the best place to go is near the current best particle. The second part, due to the diffusion term, encourages each particle to explore the space. Clearly this is exploring behaviour, and the diffusion coefficient means that more exploration is encouraged when the particle is far from the consensus. Or, equivalently, when the particles get close to the consensus, their exploring behaviour vanishes.

It would be reasonable to question why the noise in (4.9) needs to be component-wise. In other words, could a simpler model use isotropic noise:

$$dX^i(t) = -\beta(X^i(t) - \bar{X}(t))dt + \sigma|X^i(t) - \bar{X}(t)|dW^i(t). \quad (4.11)$$

In fact, noise of this form was proposed originally in [111] and then later considered in [28]. The problem, however, is that the hyperparameters, β and σ , must become dependent on the dimension d in order for convergence of the particles. This is illustrated in [30] (see also [71, Section 2.1]). It is for this reason we consider models like (4.9).

This smooth approximation to the minimum is formally justified by

the **Laplace principle**, a classical result in large deviations theory [37]. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be continuous and let μ be a probability measure with full support. Then,

$$\lim_{\alpha \rightarrow \infty} -\frac{1}{\alpha} \log \int_{\mathbb{R}^d} e^{-\alpha f(x)} \mu(\mathrm{d}x) = \inf_{x \in \mathbb{R}^d} f(x). \quad (4.12)$$

As $\alpha \rightarrow \infty$, the measure $e^{-\alpha f(x)} \mu(\mathrm{d}x)$ concentrates around the global minimizers of f , and the integral is dominated by contributions near these points.

CBO has been further developed in [28, 30, 61, 71]. A survey on the topic is available in [134].

Constrained consensus-based optimization

We provide a brief description of existing work on constrained consensus-based optimization (CCBO) [18, 6, 31].

In [18], a penalty function is added to the objective function f , which penalizes the objective when outside the feasible set (f is defined on \mathbb{R}^d). For a penalization multiplier $\lambda > 0$, the following unconstrained problem is considered:

$$\min_{x \in \mathbb{R}^d} P(x; \lambda), \quad (4.13)$$

where $P(x; \lambda) = f(x) + \lambda r(x)$. Here, $r(x)$ is a penalty function of the form $r(x) = 0$ if $x \in \bar{G}$ and $r(x) > 0$ otherwise. The particle system

$$\mathrm{d}X^i = -\beta(X^i - \bar{X})\mathrm{d}t + \sigma|X^i - \bar{X}|\mathrm{d}W^i, \quad (4.14)$$

$i = 1, \dots, N$, is considered, where

$$\bar{X} := \bar{X}^{P,\alpha} := \frac{\sum_{i=1}^N X^i e^{-\alpha P(X^i; \lambda)}}{\sum_{i=1}^N e^{-\alpha P(X^i; \lambda)}}, \quad (4.15)$$

is the particles' consensus and W^i , $i = 1, \dots, N$ are N independent standard Wiener processes. The particles are not constrained to the feasible set \bar{G} and may take any value in \mathbb{R}^d . An Euler-Maruyama scheme is applied to simulate the particle system. Convergence of the algorithm to

the optimal solution is shown, both when the exact penalty parameter λ is known, and also when it is iteratively updated.

In [6], a projection scheme is considered. Particles are updated according to an Euler discretization of the system

$$dX^i = -\beta(X^i - \bar{X})dt + \sigma(X^i - \bar{X})dW^i, \quad (4.16)$$

$i = 1, \dots, N$, where $W^i, i = 1, \dots, N$ are N d -dimensional independent standard Wiener processes. However, at each step they are projected onto the feasible space \bar{G} . No continuous-time analogue is considered; however, convergence of the iterates to a consensus state is shown and, under some mild assumptions on the objective function, this consensus state can minimize the objective function arbitrarily well.

In [31] a model of the form (4.16) is also considered; however, like [18], they also use a penalized objective function. Several numerical experiments are presented, although there are no convergence results.

Finally, we mention the recent paper [8], which was developed independently, but at the same time, as our work [65]. Similar to our work, they use reflected McKean-Vlasov SDEs to prove convergence of a CBO algorithm to the global minimizer of a constrained optimization problem. They also consider heuristics for cases where there are only a small number of particles. This is practically useful, since in many computationally demanding optimization problems (e.g. in high dimensions, or with an expensive objective function) only a few particles may be computationally feasible.

4.3 Reflected McKean-Vlasov stochastic differential equations

In this section, we introduce reflected McKean-Vlasov SDEs and their particle approximations. We then consider their application in solving constrained optimization problems.

Let $G \subset \mathbb{R}^d$ denote a bounded domain, to which our processes will

be constrained, and let $\nu(x)$ denote the unit inward normal at a point x belonging to the boundary ∂G . We assume for now that the boundary is smooth enough such that this inward normal is unique. Let $\mathcal{P}(\bar{G})$ be the space of probability measures on \bar{G} and $b : \mathbb{R}_+ \times \bar{G} \times \mathcal{P}(\bar{G}) \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}_+ \times \bar{G} \times \mathcal{P}(\bar{G}) \rightarrow \mathbb{R}^{d \times d}$. We will impose conditions on both G and the coefficients, b and σ , later.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete, sufficiently rich probability space and $(\mathcal{F}_t, 0 \leq t \leq T)$, be a filtration satisfying the usual hypothesis. Let $(W(t), \mathcal{F}_t)$ and $(\mathbf{W}(t), \mathcal{F}_t)$ be standard d -dimensional and Nd -dimensional Wiener processes, respectively. Here, $N \in \mathbb{N}$ will represent the number of particles in the considered system. We write $\mathbf{W}(t) = (W^1(t)^\top, \dots, W^N(t)^\top)^\top$. We will also use the notation: \mathcal{F}_t^W is the natural filtration for the Wiener process $W(t)$ and $\mathcal{F}_t^{\mathbf{W}}$ is the natural filtration for the Wiener process $\mathbf{W}(t)$.

We consider the non-linear (in the sense of McKean) Markov process X which satisfies the SDE:

$$\begin{aligned} X(t) = X(0) &+ \int_0^t b(s, X(s), \mathcal{L}_{X(s)}) ds + \int_0^t \sigma(s, X(s), \mathcal{L}_{X(s)}) dW(s) \\ &+ \int_0^t \nu(X(s)) dL(s), \quad X(0) \in \bar{G}, \end{aligned} \quad (4.17)$$

where $\mathcal{L}_{X(t)}$ denotes the law of $X(t)$ and $L(t)$ is a scalar non-decreasing process which only increases when $X(t) \in \partial G$. The initial condition $X(0)$ can be random but must be \mathcal{F}_0 -measurable.

To be completely precise, given a domain G and coefficients b and σ we say that the reflected McKean-Vlasov SDE (4.17) is **well-posed** if there exists a unique pair of \mathcal{F}_t^W -adapted processes (X, L) such that

1. The integral equation (4.17) holds with all integrals well-defined.
2. $X(t) \in \bar{G}$ for all $t \geq 0$.
- 3.

$$L(t) = \int_0^t I_{\partial G}(X(s)) dL(s). \quad (4.18)$$

For approximation of (4.17), we consider the particle approximation:

$$\begin{aligned}
X^{i,N}(t) &= X^{i,N}(0) + \int_0^t b\left(s, X^{i,N}(s), \hat{\mu}_{\mathbf{X}^N(s)}\right) ds \\
&+ \int_0^t \sigma\left(s, X^{i,N}(s), \hat{\mu}_{\mathbf{X}^N(s)}\right) dW^i(s) \\
&+ \int_0^t \nu(X^{i,N}(s)) dL^{i,N}(s), \quad X^{i,N}(0) \in \bar{G}, \quad (4.19)
\end{aligned}$$

where $\mathbf{X}^N(t) = (X^{1,N}(t), \dots, X^{N,N}(t))^\top$ and $L^{i,N}(t)$ is the local time of the i -th particle on the boundary ∂G . As before, $\hat{\mu}_{\mathbf{X}^N(t)}$ denotes the empirical measure of $\mathbf{X}^N(t)$:

$$\hat{\mu}_{\mathbf{X}^N(t)} = \frac{1}{N} \sum_{i=1}^N \delta_{X^{i,N}(t)}. \quad (4.20)$$

The initial conditions $X^{i,N}(0)$ are independent copies of $X(0)$.

Well-posedness of the particle system is defined analogously to well-posedness of the McKean-Vlasov SDE: we require the existence and uniqueness of $(X^{i,N}, L^{i,N})$, $i = 1, \dots, N$ such that

1. The integral equations (4.19) holds with all integrals well-defined.
2. $X^{i,N}(t) \in \bar{G}$ for all $t \geq 0$ for each $i = 1, \dots, N$.
3. Each $L^{i,N}$ is a scalar non-decreasing process satisfying

$$L^{i,N}(t) = \int_0^t I_{\partial G}(X^{i,N}(s)) dL^{i,N}(s). \quad (4.21)$$

The McKean-Vlasov SDE (4.17) and its corresponding particle system (4.19) are general and we prove well-posedness for them in Sections 4.4.1 and 4.4.2. Different choices for b and σ leads to different models. Here, we consider two models for constrained optimization.

4.3.1 Reflected consensus-based models for optimization

In this section, we consider consensus-based models in the form of reflected SDEs which can be used for solving the constrained optimization problem:

$$\min_{x \in \bar{G}} f(x), \quad (4.22)$$

where $f : \bar{G} \rightarrow [0, \infty)$ is an objective function (which can be non-convex). Since \bar{G} is compact, f obtains its infimum and supremum over \bar{G} .

Assumption 4.1 (Objective function, unique minimizer). *The objective function satisfies*

$$f_{\min} := \inf_{x \in \bar{G}} f(x) > 0, \quad (4.23)$$

and the minimizer of f , denoted as x_{\min} is unique, i.e., there is only $x_{\min} \in \bar{G}$ such that $f(x_{\min}) = f_{\min}$.

We consider the following dynamics

$$\begin{aligned} dX(t) &= -\beta(X(t) - \bar{X}(t))dt + \sigma \text{Diag}(X(t) - \bar{X}(t))dW(t) \\ &+ \nu(X(t))dL(t), \quad X(0) = X_0, \end{aligned} \quad (4.24)$$

where the weighted mean $\bar{X}(t)$ (depending on the objective f and some $\alpha > 0$) is

$$\bar{X}(t) := \frac{\int_{\bar{G}} x e^{-\alpha f(x)} \mathcal{L}_{X(t)}(dx)}{\int_{\bar{G}} e^{-\alpha f(x)} \mathcal{L}_{X(t)}(dx)}, \quad (4.25)$$

$\beta > 0$ and $\sigma > 0$ are constants (see a discussion in [71] on potential benefits of choosing them to be dependent on time). The distribution of the initial condition X_0 is assumed to have full support on \bar{G} . The relationship between β and σ required for the convergence of $X(t)$ to the global minimizer x_{\min} will be established within the convergence analysis in Section 4.6.

The particle approximation of (4.24) is given by

$$\begin{aligned} dX^{i,N}(t) &= -\beta(X^{i,N}(t) - \bar{X}^N(t))dt + \sigma \text{Diag}(X^{i,N}(t) - \bar{X}^N(t))dW^i(t) \\ &+ \nu(X^{i,N}(t))dL^{i,N}(t), \quad X^{i,N}(0) = X_0^{i,N}, \end{aligned} \quad (4.26)$$

where the weighted mean $\bar{X}^N(t)$ is

$$\bar{X}^N(t) := \bar{X}^{N,f,\alpha}(t) := \frac{\sum_{i=1}^N X^{i,N}(t) e^{-\alpha f(X^{i,N}(t))}}{\sum_{i=1}^N e^{-\alpha f(X^{i,N}(t))}}, \quad (4.27)$$

which we will refer to as the particles' consensus. In the corresponding particle system (4.26), the particles interact with each other via (4.27) and try to realize a uniform consensus.

In the reflected CBO model (4.26), the particles are attracted towards their weighted mean, which is exploitative behaviour based on already searched space. Meanwhile, the exploration is achieved thanks to noise induced by independent Brownian motions driving each particle. In [71], the exploration is enhanced by adding compound Poisson processes. Another way to facilitate exploration is by incorporating repelling forces among particles with a decaying parameter. A model incorporating both attractive and repelling behaviour takes the form

$$\begin{aligned} dX(t) = & -\beta(X(t) - \bar{X}(t))dt \\ & + \lambda(t) \int_{\bar{G}} (X(t) - y) \exp\left(-\frac{1}{2}|X(t) - y|^2\right) \mathcal{L}_{X(t)}(dy) dt \\ & + \sigma \text{Diag}(X(t) - \bar{X}(t))dW(t) + \nu(X(t))dL(t), \end{aligned} \quad (4.28)$$

and its particle approximation is given by

$$\begin{aligned} dX^{i,N}(t) = & -\beta(X^{i,N}(t) - \bar{X}^N(t))dt \\ & + \frac{\lambda(t)}{N} \sum_{j=1}^N (X^{i,N}(t) - X^{j,N}(t)) \exp\left(-\frac{1}{2}|X^{i,N}(t) - X^{j,N}(t)|^2\right) dt \\ & + \sigma \text{Diag}((X^{i,N}(t) - \bar{X}^N(t)))dW^i(t) + \nu(X^{i,N}(t))dL^{i,N}(t). \end{aligned} \quad (4.29)$$

Here $\beta > 0$ and $\sigma > 0$ are constants, and $\lambda(t) \geq 0$ is a decreasing function of t and preferably should have exponential decay in later steps of the method. The repelling force between any two particles decays as the distance between the two particles increases, in order to ensure there is no explosion in the dynamics. When particles are close to each other, they experience greater repelling to avoid collapse of the ensemble at a local minimum. The relation among β , σ and λ will be discussed in Section 4.6.

For analysis in Section 4.4, we make the following assumption on the objective problem.

Assumption 4.2 (Lipschitz objective function). *There exists $L_f > 0$ such that*

$$|f(x) - f(y)| \leq L_f |x - y|, \quad (4.30)$$

for all $x, y \in \bar{G}$.

Assumption 4.2 ensures the coefficients of the McKean-Vlasov SDE and its particle approximations are indeed Lipschitz. If f is continuously differentiable then (4.30) is automatically satisfied. We denote the supremum of f as

$$f_{\max} = \sup_{x \in \bar{G}} f(x). \quad (4.31)$$

4.4 Well-posedness results

In this section, we consider the general reflected McKean-Vlasov SDE (4.17) and its particle approximation (4.19). Our aim is to show that these SDEs are well-posed; that is, for both the mean-field limit and the particle approximation, there exists a unique process which does indeed satisfy the corresponding integral equations. We prove well-posedness of the mean-field SDE in Section 4.4.1 and well-posedness of the particle system in Section 4.4.2. Certain conditions are imposed on the coefficients in order to show well-posedness. We verify that the models considered in Section 4.3 satisfy these conditions in Section 4.4.3. We begin this section by recalling fixed-point arguments, which we make use of in the well-posedness results.

Contraction-based fixed point arguments

The standard way to prove existence and uniqueness of the reflected McKean-Vlasov SDE (4.17) and its particle approximation (4.19) is to appeal to a contraction-based fixed point argument, see for instance the method in [131, 90] or more recently [1]. We briefly recall the Banach fixed-point theorem and explain how it is applied to establish existence and uniqueness.

Let $(V, \|\cdot\|)$ be a (non-empty) complete normed vector space (a Banach space). A map $\Phi : V \rightarrow V$ is a contraction if there exists a constant $q \in [0, 1)$ such that

$$\|\Phi(x) - \Phi(y)\| \leq q\|x - y\|,$$

for any $x, y \in V$. The Banach fixed-point theorem states that any contraction Φ admits a unique fixed point x^* , i.e. there exists an $x^* \in V$ such that $\Phi(x^*) = x^*$. To see this, start with any $x_0 \in V$, then construct a sequence by iteratively applying Φ , so that $x_{k+1} = \Phi(x_k)$. The sequence $\{x_k\}$ is a Cauchy sequence due to the contraction property of Φ . Completeness of the V yields the result [83].

As a starting example, consider the initial-value problem for the ordinary differential equation

$$\frac{dy}{dt} = f(t, y(t)), \quad y(t_0) = y_0. \quad (4.32)$$

Taking $V = C([t_0, t_0+T]; \mathbb{R}^d)$, for some $T > 0$, with $\|y\| = \sup_{s \in [t_0, t_0+T]} |y(s)|$, consider the integral operator $\Phi : V \rightarrow V$ defined by

$$\Phi(y)(t) = y_0 + \int_{t_0}^t f(s, y(s)) ds, \quad (4.33)$$

$t \in [t_0, t_0 + T]$. Note that if there exists some $y \in V$ with $y = \Phi(y)$ the integral equation related to (4.32) is satisfied. Thus, establishing existence and uniqueness of (4.32) amounts to showing that Φ has a fixed point. This is done by making assumptions on f such that Φ becomes a contraction. For ordinary differential equations, this is called the Picard-Lindelöf theorem [143] and a similar argument is used to prove existence and uniqueness of SDEs [108].

Finally, we note that we do not necessarily need Φ to be a contraction in order for it to admit a fixed point. A weaker condition, which we make use of in the sequel, is that the j -fold composition of Φ (denoted by Φ^j) is a contraction. To see why this condition implies that Φ has a fixed point, note that Φ^j has a fixed point by the Banach fixed point theorem, which

we can denote x^* . Then,

$$\begin{aligned}\Phi^j(\Phi(x^*)) &= \Phi(\Phi^j(x^*)) \\ &= \Phi(x^*)\end{aligned}$$

and hence $\Phi(x^*)$ must also be a fixed point of Φ^j . By uniqueness of the fixed point of Φ^j we must have $\Phi(x^*) = x^*$, so x^* is a fixed point of Φ also [60].

4.4.1 Well-posedness of reflected McKean-Vlasov SDEs

We now prove well-posedness of the reflected McKean-Vlasov SDE (4.17). When applying the fixed-point argument for reflected SDEs, a key consideration is how to deal with the local time terms. If G were a convex domain, then it is easier to deal with the local time term in (4.17) using Tanaka's trick [132] which is exploited in the McKean-Vlasov setting in [1]. Here we do not assume that G is convex, but instead assume that G is bounded and rely on the boundary ∂G satisfying the uniform exterior sphere condition. Sznitman [131] considered this case of non-convex G but only for a first-order interaction. Note that the optimization models discussed in Section 4.3.1 do not fall into this category, and hence we study the general mean-field SDEs (4.17). We borrow Sznitman's [131] idea on how to deal with the local time terms and apply it within our setting.

For a measurable space $(S, \mathcal{B}(S))$, denote by $\mathcal{P}(S)$ the space of probability measures on S . We use the notation: $\mathcal{C} = C(\mathbb{R}_+, \bar{G})$ equipped with the Borel σ -algebra generated by the uniform norm topology and $\mathcal{M} = \mathcal{P}(\mathcal{C})$. For a measure $\mu \in \mathcal{M}$, we denote by μ_t its projection onto $\mathcal{P}(\bar{G})$ at time t , i.e. its marginal:

$$\mu_t(A) = \int_{\mathcal{C}} I_A(w_t) d\mu(w), \quad A \in \mathcal{B}(\bar{G}). \quad (4.34)$$

For two measures $\mu^1, \mu^2 \in \mathcal{P}(\bar{G})$, the Wasserstein p -metric is defined as

$$\mathcal{W}_p(\mu^1, \mu^2) := \inf_{\gamma \in \Gamma(\mu^1, \mu^2)} \left[\int_{\bar{G} \times \bar{G}} |x - y|^p d\gamma(x, y) \right]^{\frac{1}{p}}, \quad (4.35)$$

where $\Gamma(\mu^1, \mu^2)$ denotes the set of couplings between μ^1 and μ^2 . One important property of the Wasserstein p -metric is that the distance is monotone increasing in the order p . For two measures μ^1, μ^2 and for $1 \leq p \leq q < \infty$ we have

$$\mathcal{W}_p(\mu^1, \mu^2) \leq \mathcal{W}_q(\mu^1, \mu^2), \quad (4.36)$$

see for instance [140].

Assumption 4.3. *The domain $G \subset \mathbb{R}^d$ is bounded and its boundary ∂G is C^3 .*

The smoothness of the boundary assumed here is required so that the distance function to the boundary of G , $d(\cdot, \partial G)$, defined on a neighbourhood of ∂G is C^2 (see e.g. [120, Section 3, Lemma 1]). We make use of this in Theorem 4.1.

Assumption 4.4 (Uniformly Lipschitz in space and measure). *There exists $L > 0$ such that for all $t \geq 0$*

$$|b(t, x, \mu^1) - b(t, y, \mu^2)| + |\sigma(t, x, \mu^1) - \sigma(t, y, \mu^2)| \leq L(|x - y| + \mathcal{W}_4(\mu^1, \mu^2)), \\ \forall x, y \in \bar{G}, \forall \mu^1, \mu^2 \in \mathcal{P}(\bar{G}). \quad (4.37)$$

Also, the coefficients $b(t, x, \mu)$ and $\sigma(t, x, \mu)$ are continuous in t .

Since we are working on a bounded domain, the use of the \mathcal{W}_4 (and hence requiring fourth moments of the measures) is not more restrictive than using \mathcal{W}_2 .

Since ∂G is sufficiently smooth, it satisfies the uniform exterior sphere condition. That is, there exists $R_0 > 0$ such that $\forall x \in \partial G$,

$$\bar{B}(x - R_0\nu(x), R_0) \cap \partial G = \{x\},$$

where $\bar{B}(x, R_0)$ denotes the closed ball of radius R_0 centred at x with respect to the Euclidean norm, and $\nu(x)$ is the inward normal vector field at $x \in \partial G$. The constant R_0 is called the uniform exterior sphere constant of G . This is equivalent to the following condition:

$$\exists c \geq 0, \quad \forall x \in \partial G, \quad \forall y \in \bar{G}, \quad c|x - y|^2 \geq (x - y)^\top \nu(x), \quad (4.38)$$

see for example [110]. The constant c is related to the uniform exterior sphere constant by $c = \frac{1}{2R_0}$. In the convex case we have $c = 0$.

Theorem 4.1. *Let Assumptions 4.3 and 4.4 hold. There exists a unique pair of continuous \mathcal{F}_t^W -adapted processes $(X(t), L(t))$ such that*

1. $X(t) \in \bar{G}$ for all $t \geq 0$
2. $L(t)$ is non-decreasing with $L(0) = 0$ and for all $t \geq 0$,

$$L(t) = \int_0^t I_{\partial G}(X(s)) dL(s) \quad (4.39)$$

3. For all $t \geq 0$,

$$\begin{aligned} X(t) = X(0) &+ \int_0^t b(s, X(s), \mathcal{L}_{X(s)}) ds + \int_0^t \sigma(s, X(s), \mathcal{L}_{X(s)}) dW(s) \\ &+ \int_0^t \nu(X(s)) dL(s). \end{aligned} \quad (4.40)$$

Such a theorem is proved in [131] in the case of first-order interaction rather than (4.17). Our proof of Theorem 4.1 borrows certain arguments and notation from [131] which also draws techniques from [90] to deal with the local time term.

For brevity, we denote $X_s := X(s)$, $W_s := W(s)$, and $L_s := L(s)$. Consider the map $\Phi : \mathcal{M} \rightarrow \mathcal{M}$ defined by $\Phi(\mu) = \mathcal{L}_{Y^\mu}$, where \mathcal{L}_{Y^μ} is the law of $\{Y_t^\mu\}_{t \geq 0}$ which is defined as

$$Y_t^\mu = X_0 + \int_0^t b(s, Y_s^\mu, \mu_s) ds + \int_0^t \sigma(s, Y_s^\mu, \mu_s) dW_s + \int_0^t \nu(Y_s^\mu) dL_s. \quad (4.41)$$

The process Y is decoupled from its own measure and, as such, is just

an ordinary reflected SDE. The existence and uniqueness of a solution to (4.41) is proved in [90] when the coefficients b and σ are defined on the whole space \mathbb{R}^d . We note that we can extend our coefficients smoothly to the whole space, and the choice of extension does not affect the solution Y^μ . Moreover, the assumptions required in [90] are all satisfied in our setting. Hence, the map Φ is well-defined. From the definition of Φ , if Φ has a unique fixed point, then the processes associated with this fixed point are the unique solution to (4.17). Our objective is to show that there exists a unique fixed point by demonstrating that, for some j , the j -fold composition, Φ^j , is a contraction.

Let us fix an arbitrary $T > 0$ and consider the space $\mathcal{C}_T = C([0, T], \bar{G})$ and $\mathcal{M}_T := \mathcal{P}(\mathcal{C}_T)$ equipped with the metric

$$D_T(\mu^1, \mu^2) = \inf_{\gamma \in \Gamma(\mu^1, \mu^2)} \left[\int_{\mathcal{C}_T \times \mathcal{C}_T} \left(\sup_{s \leq T} |X(s) - Y(s)| \right)^4 d\gamma(X, Y) \right]^{\frac{1}{4}}.$$

Note that D_T is the Wasserstein 4-distance with respect to the norm $\sup_{s \leq T} |X(s)|$, $X \in \mathcal{C}_T$. Completeness of the space (\mathcal{M}_T, D_T) follows from separability and completeness of $(\mathcal{C}_T, \sup_{s \leq T} |X(s) - Y(s)|)$ (see [16, 17]). Note that although D_T is a metric on \mathcal{M}_T , we will abuse notation and write $D_T(\mu^1, \mu^2)$ for $\mu^1, \mu^2 \in \mathcal{M}$, and by this we mean the D_T -distance of the projections of μ^1, μ^2 onto \mathcal{M}_T .

To prove Theorem 4.1, we need the following lemma.

Lemma 4.2. *Under Assumptions 4.3 and 4.4, there exists $K_T > 0$ such that for all $\mu_1, \mu_2 \in \mathcal{M}$,*

$$D_T^4(\Phi(\mu^1), \Phi(\mu^2)) \leq K_T \int_0^T D_u^4(\mu^1, \mu^2) du. \quad (4.42)$$

Proof. Let $\mu^1, \mu^2 \in \mathcal{M}$ and define the processes $Y_t^1 := Y_t^{\mu^1}$, $Y_t^2 := Y_t^{\mu^2}$ via (4.41). Denote by $d(\cdot, \partial G)$ the Euclidean distance to the boundary of G , defined on some neighborhood of ∂G within G and let the function $g(\cdot)$ be a smooth, bounded extension of it to the whole space. The choice of

extension does not matter but we will rely on the fact that $\nabla g(x) = \nu(x)$ when $x \in \partial G$. Smoothness and boundedness of g will also be used in the proof. Smoothness of $d(\cdot, \partial G)$ is not guaranteed on all of G , only sufficiently close enough to the boundary, due to the closest boundary point potentially switching. Smoothness of g is required for application of Ito's formula. Then Ito's formula gives

$$\begin{aligned}
g(Y_t^i) &= g(X_0) + \int_0^t \nabla g^\top(Y_s^i) b(s, Y_s^i, \mu_s^i) ds \\
&\quad + \int_0^t \frac{1}{2} \text{tr}[\sigma(s, Y_s^i, \mu_s^i)^\top \nabla^2 g(Y_s^i) \sigma(s, Y_s^i, \mu_s^i)] ds \\
&\quad + \int_0^t \nabla g^\top(Y_s^i) \sigma(s, Y_s^i, \mu_s^i) dW_s + \int_0^t \nabla g^\top(Y_s^i) \nu(Y_s^i) dL_s
\end{aligned} \tag{4.43}$$

for $i = 1, 2$, where ∇^2 denotes the Hessian operator. Noting that $\nabla g = \nu$ on ∂G , we have

$$\int_0^t \nabla g^\top(Y_s^i) \nu(Y_s^i) dL_s = L_t. \tag{4.44}$$

In the interest of brevity, we write this as

$$g_t^i = g_0 + \int_0^t \left(\nabla g_s^{i\top} b_s^i + \frac{1}{2} \text{tr}[\sigma_s^{i\top} \nabla^2 g_s^i \sigma_s^i] \right) ds + \int_0^t \nabla g_s^{i\top} \sigma_s^i dW_s + L_t^i, \quad i = 1, 2. \tag{4.45}$$

Let $c > 0$ be the uniform exterior sphere constant of G from (4.38).

By Ito's formula, we have

$$\begin{aligned}
\exp\{-2c(g_t^1 + g_t^2)\} &= \exp\{-4cg_0\} \\
&\quad + \int_0^t \exp\{-2c(g_s^1 + g_s^2)\} \left(-2c\tilde{b}_s + 2c^2\tilde{\sigma}_s\tilde{\sigma}_s^\top \right) ds \\
&\quad - \int_0^t 2c \exp\{-2c(g_s^1 + g_s^2)\} \tilde{\sigma} dW_s \\
&\quad - \int_0^t 2c \exp\{-2c(g_s^1 + g_s^2)\} [dL_s^1 + dL_s^2], \tag{4.46}
\end{aligned}$$

where

$$\begin{aligned}\tilde{b}_s &= (\nabla g_s^1)^\top b_s^1 + (\nabla g_s^2)^\top b_s^2 + \frac{1}{2} \text{tr}[(\sigma_s^1)^\top \nabla^2 g_s^1 \sigma_s^1 + (\sigma_s^2)^\top \nabla^2 g_s^2 \sigma_s^2], \\ \tilde{\sigma}_s &= (\nabla g_s^1)^\top \sigma_s^1 + (\nabla g_s^2)^\top \sigma_s^2.\end{aligned}\tag{4.47}$$

Also,

$$\begin{aligned}|Y_t^1 - Y_t^2|^2 &= \int_0^t 2(Y_s^1 - Y_s^2)^\top (b_s^1 - b_s^2) + \text{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)] \mathbf{d}s \\ &+ \int_0^t 2(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \mathbf{d}W_s + \int_0^t 2(Y_s^1 - Y_s^2)^\top [\nu(Y_s^1) \mathbf{d}L_s^1 - \nu(Y_s^2) \mathbf{d}L_s^2].\end{aligned}\tag{4.48}$$

Then, for the product, Ito's formula yields

$$\begin{aligned}&\exp\{-2c(g_t^1 + g_t^2)\} |Y_t^1 - Y_t^2|^2 \\ &= 2 \int_0^t \exp\{-2c(g_s^1 + g_s^2)\} (Y_s^1 - Y_s^2)^\top [(b_s^1 - b_s^2) \mathbf{d}s + (\sigma_s^1 - \sigma_s^2) \mathbf{d}W_s \\ &\quad + \nu(Y_s^1) \mathbf{d}L_s^1 - \nu(Y_s^2) \mathbf{d}L_s^2] \\ &\quad + \int_0^t \exp\{-2c(g_s^1 + g_s^2)\} \text{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)] \mathbf{d}s \\ &\quad - 2c \int_0^t |Y_s^1 - Y_s^2|^2 \exp\{-2c(g_s^1 + g_s^2)\} [\tilde{b}_s \mathbf{d}s + \tilde{\sigma}_s^\top \mathbf{d}W_s + \mathbf{d}L_s^1 + \mathbf{d}L_s^2] \\ &\quad + 2c^2 \int_0^t \exp\{-2c(g_s^1 + g_s^2)\} |Y_s^1 - Y_s^2|^2 \tilde{\sigma}_s \tilde{\sigma}_s^\top \mathbf{d}s \\ &\quad - 4c \int_0^t \exp\{-2c(g_s^1 + g_s^2)\} (Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top \mathbf{d}s.\end{aligned}\tag{4.49}$$

From the uniform exterior sphere condition (4.38), we have

$$\begin{aligned}-c|Y_s^1 - Y_s^2|^2 + (Y_s^1 - Y_s^2)^\top \nu(Y_s^1) &\leq 0, \\ -c|Y_s^1 - Y_s^2|^2 + (Y_s^2 - Y_s^1)^\top \nu(Y_s^2) &\leq 0, \text{ a.s.}\end{aligned}\tag{4.50}$$

To ease the notation, we denote $\kappa(s) := \exp\{-2c(g_s^1 + g_s^2)\}$; note that κ is

bounded since g is bounded. Then,

$$\begin{aligned}
\kappa(t)|Y_t^1 - Y_t^2|^2 &\leq 2 \int_0^t \kappa(s)(Y_s^1 - Y_s^2)^\top [(b_s^1 - b_s^2)\mathbf{d}s + (\sigma_s^1 - \sigma_s^2)\mathbf{d}W_s] \\
&+ \int_0^t \kappa(s)\mathbf{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)]\mathbf{d}s - 2c \int_0^t |Y_s^1 - Y_s^2|^2 \kappa(s)[\tilde{b}_s\mathbf{d}s + \tilde{\sigma}_s^\top \mathbf{d}W_s] \\
&+ 2c^2 \int_0^t \kappa(s)|Y_s^1 - Y_s^2|^2 \tilde{\sigma}_s \tilde{\sigma}_s^\top \mathbf{d}s - 4c \int_0^t \kappa(s)(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top \mathbf{d}s.
\end{aligned} \tag{4.51}$$

Squaring both sides and applying the Cauchy-Schwarz inequality twice, we obtain for some constant $K > 0$

$$\begin{aligned}
|Y_t^1 - Y_t^2|^4 &\leq K \left[\int_0^t [(Y_s^1 - Y_s^2)^\top (b_s^1 - b_s^2)]^2 \mathbf{d}s \right. \\
&+ \int_0^t [\mathbf{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)]]^2 \mathbf{d}s \\
&+ \int_0^t |Y_s^1 - Y_s^2|^4 \mathbf{d}s + \int_0^t [(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top]^2 \mathbf{d}s \\
&\left. + \left(\int_0^t \kappa(s)(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \mathbf{d}W_s \right)^2 + \left(\int_0^t \kappa(s)|Y_s^1 - Y_s^2|^2 \tilde{\sigma}_s \mathbf{d}W_s \right)^2 \right],
\end{aligned} \tag{4.52}$$

where we use the fact that $\kappa(s)$ is bounded by definition, and \tilde{b} and $\tilde{\sigma}$ are bounded on \bar{G} . This boundedness follows from the fact that $\nabla g, \nabla^2 g, b,$ and σ are all continuous and therefore bounded on the compact set \bar{G} . Then, taking the supremum over $[0, t]$ and taking expectation, we get

$$\begin{aligned}
\mathbb{E} \sup_{u \leq t} |Y_u^1 - Y_u^2|^4 &\leq K \left[\mathbb{E} \sup_{u \leq t} \left(\int_0^u [(Y_s^1 - Y_s^2)^\top (b_s^1 - b_s^2)]^2 \mathbf{d}s \right. \right. \\
&+ \int_0^u [\mathbf{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)]]^2 \mathbf{d}s \\
&+ \int_0^u |Y_s^1 - Y_s^2|^4 \mathbf{d}s + \int_0^u [(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top]^2 \mathbf{d}s \Big) \\
&+ \mathbb{E} \sup_{u \leq t} \left(\int_0^u \kappa(s)(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \mathbf{d}W_s \right)^2 \\
&\left. + \mathbb{E} \sup_{u \leq t} \left(\int_0^u \kappa(s)|Y_s^1 - Y_s^2|^2 \tilde{\sigma}_s \mathbf{d}W_s \right)^2 \right].
\end{aligned} \tag{4.53}$$

The Riemann integrals have non-negative integrands so their supremum is reached at t . Meanwhile for the Ito integrals, we use Doob's maximal inequality followed by the Ito isometry, which yields

$$\begin{aligned}
\mathbb{E} \sup_{u \leq t} |Y_u^1 - Y_u^2|^4 &\leq K \mathbb{E} \left(\int_0^t [(Y_s^1 - Y_s^2)^\top (b_s^1 - b_s^2)]^2 \mathrm{d}s \right. \\
&\quad + \int_0^t [\mathrm{tr}[(\sigma_s^1 - \sigma_s^2)^\top (\sigma_s^1 - \sigma_s^2)]]^2 \mathrm{d}s \\
&\quad + \int_0^t |Y_s^1 - Y_s^2|^4 \mathrm{d}s + \int_0^t [(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top]^2 \mathrm{d}s \\
&\quad \left. + \int_0^t \mathbb{E} |(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2)|^2 \mathrm{d}s + \int_0^t \mathbb{E} |Y_s^1 - Y_s^2|^4 \mathrm{d}s \right), \tag{4.54}
\end{aligned}$$

where again we have used boundedness of $\kappa(s)$ and $\tilde{\sigma}$. Now using the Cauchy-Schwarz inequality and the Assumption 4.4, we obtain

$$\begin{aligned}
\mathbb{E} \sup_{u \leq t} |Y_u^1 - Y_u^2|^4 &\leq K \mathbb{E} \left(\int_0^t |Y_s^1 - Y_s^2|^2 |b_s^1 - b_s^2|^2 + |Y_s^1 - Y_s^2|^2 |\sigma_s^1 - \sigma_s^2|^2 \right. \\
&\quad \left. + |Y_s^1 - Y_s^2|^4 + |\sigma_s^1 - \sigma_s^2|^4 \mathrm{d}s \right) \tag{4.55}
\end{aligned}$$

$$\leq K \left(\int_0^t \mathbb{E} |Y_s^1 - Y_s^2|^4 + \mathcal{W}_4^4(\mu_s^1, \mu_s^2) \mathrm{d}s \right). \tag{4.56}$$

Note that $|(Y_s^1 - Y_s^2)^\top (\sigma_s^1 - \sigma_s^2) \tilde{\sigma}_s^\top| \leq K |Y_s^1 - Y_s^2| |\sigma_s^1 - \sigma_s^2|$ follows from application of the Cauchy-Schwarz inequality, the operator norm being bounded by Frobenius norm, and $\tilde{\sigma}$ being bounded. Then,

$$\mathbb{E} \sup_{u \leq t} |Y_u^1 - Y_u^2|^4 \leq \int_0^t K \mathbb{E} \sup_{u \leq s} |Y_u^1 - Y_u^2|^4 \mathrm{d}s + K \int_0^t \mathcal{W}_4^4(\mu_s^1, \mu_s^2) \mathrm{d}s, \tag{4.57}$$

to which Grönwall's lemma can be applied to yield

$$\mathbb{E} \sup_{u \leq t} |Y_u^1 - Y_u^2|^4 \leq K e^{Kt} \int_0^t \mathcal{W}_4^4(\mu_s^1, \mu_s^2) \mathrm{d}s. \tag{4.58}$$

Hence

$$D_T^4(\Phi(\mu^1), \Phi(\mu^2)) \leq \mathbb{E} \sup_{u \leq T} |Y_u^1 - Y_u^2|^4 \quad (4.59)$$

$$\leq K e^{KT} \int_0^T \mathcal{W}_4^4(\mu_u^1, \mu_u^2) du \quad (4.60)$$

$$\leq K e^{KT} \int_0^T D_u^4(\mu^1, \mu^2) du. \quad (4.61)$$

The inequality (4.59) follows from the definition of the metric D_T , while (4.61) is a consequence of $\mathcal{W}_4^4(\mu_u^1, \mu_u^2) \leq D_u^4(\mu^1, \mu^2)$. Lemma 4.2 is proved. \square

Now we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. Applying (4.42) twice yields

$$D_T^4(\Phi(\Phi(\mu^1)), \Phi(\Phi(\mu^2))) \leq K_T \int_0^T D_{t_1}^4(\Phi(\mu^1), \Phi(\mu^2)) dt_1 \quad (4.62)$$

$$\leq K_T^2 \int_0^T \int_0^{t_1} D_{t_2}^4(\mu^1, \mu^2) dt_2 dt_1. \quad (4.63)$$

Continuing in this way for the j -fold composition Φ^j , we obtain

$$D_T^4(\Phi^j(\mu^1), \Phi^j(\mu^2)) \leq K_T^j \int_0^T \int_0^{t_1} \cdots \int_0^{t_{j-1}} D_{t_j}^4(\mu^1, \mu^2) dt_j \cdots dt_1. \quad (4.64)$$

Changing the order of integration yields

$$D_T^4(\Phi^j(\mu^1), \Phi^j(\mu^2)) \leq K_T^j \int_0^T \int_{t_j}^T \cdots \int_{t_2}^T D_{t_j}^4(\mu^1, \mu^2) dt_1 \cdots dt_j \quad (4.65)$$

$$= K_T^j \int_0^T \frac{(T - t_j)^{j-1}}{(j-1)!} D_{t_j}^4(\mu^1, \mu^2) dt_j \leq \frac{(K_T T)^j}{j!} D_T^4(\mu^1, \mu^2). \quad (4.66)$$

Hence, if we choose j large enough such that

$$\frac{(K_T T)^j}{j!} < 1,$$

it follows that Φ^j is a contraction and Φ has a unique fixed point, as required to complete the proof. \square

4.4.2 Well-posedness of the particle system

Now we show well-posedness of the particle system (4.19).

Theorem 4.3. *Let Assumptions 4.3 and 4.4 hold. Then there exists a unique strong solution of the SDEs (4.19).*

Proof. For the sake of simplicity in presentation, we denote $X_s^{i,N} := X^{i,N}(s)$, $W_s^i := W^i(s)$ and $L_s^i := L^i(s)$. Let \bar{G}^N denote the Cartesian product of N copies of \bar{G} and let \mathbb{S} be the space of continuous and $\mathcal{F}_t^{\mathbf{W}}$ -adapted \bar{G}^N -valued processes. We define the map

$$\begin{aligned} \tilde{\Phi} : \mathbb{S} &\rightarrow \mathbb{S}, \\ Y &\mapsto \bar{Y}, \end{aligned} \tag{4.67}$$

where $\bar{Y} = (\bar{Y}^{1,N^\top}, \dots, \bar{Y}^{N,N^\top})^\top$ is the solution to

$$\begin{aligned} \bar{Y}_t^{i,N} &= X_0^{i,N} + \int_0^t b(s, \bar{Y}_s^{i,N}, \hat{\mu}_{Y_s}) \mathrm{d}s + \int_0^t \sigma(s, \bar{Y}_s^{i,N}, \hat{\mu}_{Y_s}) \mathrm{d}W_s^i \\ &\quad + \int_0^t \nu(\bar{Y}_s^{i,N}) \mathrm{d}L_s^{i,N}, \end{aligned} \tag{4.68}$$

$i = 1, \dots, N$, and $\hat{\mu}_{Y_t}$ is the empirical measure of an input process $Y \in \mathbb{S}$ at time t (see (4.20)). The SDE (4.68) is a reflected SDE with random coefficients. Existence and uniqueness of each component \bar{Y}^i follow from [21, Theorem 2.2]. Hence, the map $\tilde{\Phi}$ is well-defined.

Let $Y^{(1)}, Y^{(2)} \in \mathbb{S}$ and $\bar{Y}^{(j)} = \tilde{\Phi}(Y^{(j)})$, $j = 1, 2$. Applying the technique of Lemma 4.2 to each of the N components of the two processes $\bar{Y}^{(j)}$, $j = 1, 2$, we obtain (cf. (4.56))

$$\begin{aligned} \mathbb{E} \sup_{u \leq t} |\bar{Y}_u^{(1),i,N} - \bar{Y}_u^{(2),i,N}|^4 &\leq K \left(\int_0^t \mathbb{E} |\bar{Y}_s^{(1),i,N} - \bar{Y}_s^{(2),i,N}|^4 \right. \\ &\quad \left. + \mathbb{E} \mathcal{W}_4^4 \left(\hat{\mu}_{Y_s^{(1)}}, \hat{\mu}_{Y_s^{(2)}} \right) \mathrm{d}s \right), \end{aligned} \tag{4.69}$$

for $i = 1, \dots, N$. Note that in comparison to (4.56), we now have an expectation of the Wasserstein distance due to the fact that the measures

are empirical. Applying Grönwall's lemma yields

$$\mathbb{E} \sup_{u \leq t} |\bar{Y}_u^{(1),i,N} - \bar{Y}_u^{(2),i,N}|^4 \leq K e^{Kt} \left(\int_0^t \mathbb{E} \mathcal{W}_4^4(\hat{\mu}_{Y_s^{(1)}}, \hat{\mu}_{Y_s^{(2)}}) \mathbf{d}s \right). \quad (4.70)$$

From (4.35) and (4.20), it can be seen that

$$\mathbb{E} \mathcal{W}_4^4(\hat{\mu}_{Y_s^{(1)}}, \hat{\mu}_{Y_s^{(2)}}) \leq \frac{1}{N} \sum_{i=1}^N \mathbb{E} |Y_s^{(1),i,N} - Y_s^{(2),i,N}|^4. \quad (4.71)$$

Hence,

$$\mathbb{E} \sup_{u \leq t} |\bar{Y}_u^{(1)} - \bar{Y}_u^{(2)}|^4 \leq K \sum_{i=1}^N \mathbb{E} \sup_{u \leq t} |\bar{Y}_u^{(1),i,N} - \bar{Y}_u^{(2),i,N}|^4 \quad (4.72)$$

$$\leq K e^{Kt} \int_0^t \sum_{i=1}^N \mathbb{E} |Y_s^{(1),i,N} - Y_s^{(2),i,N}|^4 \mathbf{d}s \quad (4.73)$$

$$\leq K e^{Kt} \int_0^t \mathbb{E} \sup_{u \leq s} |Y_u^{(1)} - Y_u^{(2)}|^4 \mathbf{d}s, \quad (4.74)$$

where for the second inequality we use (4.70) and then (4.71). One can then show that the map $\tilde{\Phi}^j$ is a contraction as before, see (4.66). The unique fixed-point of $\tilde{\Phi}^j$ is the solution to the particle system. \square

4.4.3 Well-posedness of the CBO models

We now verify that the CBO mean-field SDEs and their particle approximations from Section 4.3.1 satisfy Assumption 4.4 so that their well-posedness is established.

For $\mu \in \mathcal{P}(\bar{G})$, let us define

$$\bar{X}^\mu := \frac{\int_{\bar{G}} x e^{-\alpha f(x)} \mu(\mathbf{d}x)}{\int_{\bar{G}} e^{-\alpha f(x)} \mu(\mathbf{d}x)}. \quad (4.75)$$

The following lemma is proved in [28].

Lemma 4.4. *Let Assumption 4.3 hold and f satisfy Assumption 4.1. There exists $K > 0$ such that for all $\mu, \bar{\mu} \in \mathcal{P}(\bar{G})$,*

$$|\bar{X}^\mu - \bar{X}^{\bar{\mu}}| \leq K \mathcal{W}_2(\mu, \bar{\mu}). \quad (4.76)$$

In other words, the map $\mu \mapsto \bar{X}^\mu$ is Lipschitz with respect to the W_2 metric. Note that by monotonicity of Wasserstein distances, it is also Lipschitz with respect to W_4 . Now we show that the coefficients of the CBO model (4.24) satisfy Assumption 4.1.

Lemma 4.5. *The coefficients of (4.24),*

$$b(t, x, \mu) = \beta(x - \bar{X}^\mu) \text{ and } \sigma(t, x, \mu) = \sigma \text{Diag}(x - \bar{X}^\mu), \quad (4.77)$$

are Lipschitz, uniformly in time, with respect to the space and measure arguments (i.e., they satisfy Assumption 4.4).

Proof. For $x, \bar{x} \in \mathbb{R}^d$ and $\mu, \bar{\mu} \in \mathcal{P}(\bar{G})$,

$$\begin{aligned} |b(t, x, \mu) - b(t, \bar{x}, \bar{\mu})| + |\sigma(t, x, \mu) - \sigma(t, \bar{x}, \bar{\mu})| &= \beta|(x - \bar{x}) + (\bar{X}^{\bar{\mu}} - \bar{X}^\mu)| \\ &\quad + \sigma|(x - \bar{x}) + (\bar{X}^{\bar{\mu}} - \bar{X}^\mu)| \leq K(|x - \bar{x}| + |\bar{X}^{\bar{\mu}} - \bar{X}^\mu|) \\ &\leq K(|x - \bar{x}| + \mathcal{W}_2(\mu, \bar{\mu})) \leq K(|x - \bar{x}| + \mathcal{W}_4(\mu, \bar{\mu})), \end{aligned}$$

by application of Lemma 4.4. □

Hence, by Theorem 4.1 we have well-posedness of the mean-field limit (4.24) and the particle system (4.26). In the same way, it can be verified that the coefficients of the CBO model (4.28) satisfy Assumption 4.4, from which well-posedness of the mean-field limit (4.28) and the particle system (4.29) follows.

4.5 Convergence of interacting particle system to the mean-field limit

In this section, we prove a theorem on propagation of chaos with the optimal convergence rate. Let $X^{i,N}$, i, \dots, N , denote the solution to the particle system (4.19), where $X^{i,N}$ has driving Wiener process W^i . Let $\mathbf{X}(t) := (X^1(t), \dots, X^N(t))^\top$, where $X^i(t)$, $i = 1, \dots, N$, are i.i.d. particles solving the McKean-Vlasov SDEs (4.17) which are driven by the independent Wiener processes W^i that are the same as in the particle system

(4.19).

Under Assumption 4.3 and 4.4, it is not difficult to prove using an analogue of (4.56) and the well-known result for rates of convergence of empirical measures (see [45, Theorem 1]) that for some constant $C > 0$:

$$\max_{i=1, \dots, N} \sup_{t \leq T} \mathbb{E} |X^{i,N}(t) - X^i(t)|^4 \leq C \begin{cases} N^{-1/2}, & d < 8, \\ N^{-1/2} \log N, & d = 8, \\ N^{-\frac{4}{d}}, & d > 8. \end{cases} \quad (4.78)$$

Under the below Assumption 4.5 replacing Assumption 4.4, we instead prove the optimal convergence rate of order $1/\sqrt{N}$ for any d . The CBO models from Section 4.3.1 satisfy Assumption 4.5.

Assumption 4.5. *There exists a constant $L > 0$ such that for any $x, y \in \bar{G}$ and any $\mu_1, \mu_2 \in \mathcal{P}(\bar{G})$*

$$\begin{aligned} & |b(t, x, \mu_1) - b(t, y, \mu_2)| + |\sigma(t, x, \mu_1) - \sigma(t, y, \mu_2)| \\ & \leq L \left(|x - y| + \sum_{j=1}^J \left| \int_G \phi_j(t, x, z) d\mu_1(z) - \int_G \phi_j(t, x, z) d\mu_2(z) \right| \right), \end{aligned} \quad (4.79)$$

where $\phi_j(t, x, y)$, $j = 1, \dots, J$, are continuous functions on $[0, T] \times \bar{G} \times \bar{G}$ and for some $C > 0$

$$|\phi_j(t, x, y) - \phi_j(t, x, z)| \leq C|y - z|.$$

Also, the coefficients $b(t, x, \mu)$ and $\sigma(t, x, \mu)$ are continuous in t .

We will need the following elementary lemma.

Lemma 4.6. *Let $\{A_i\}_{i=1}^N$ be a collection of independent \mathbb{R}^d -valued random variables with uniformly bounded fourth moments and $\mathbb{E}A_i = 0$, $i = 1, \dots, N$. Then*

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N A^i \right|^4 \leq \frac{C}{N^2}, \quad (4.80)$$

where $C > 0$ is independent of N .

Proof. We have

$$\left| \sum_{i=1}^N A_i \right|^4 = \left\langle \sum_{i=1}^N A_i, \sum_{i=1}^N A_i \right\rangle^2, \quad (4.81)$$

$$= \sum_{i,j,k,\ell=1}^N \langle A_i, A_j \rangle \langle A_k, A_\ell \rangle. \quad (4.82)$$

Then

$$\mathbb{E} \left| \sum_{i=1}^N A_i \right|^4 = \sum_{i,j,k,\ell=1}^N \mathbb{E}(\langle A_i, A_j \rangle \langle A_k, A_\ell \rangle). \quad (4.83)$$

Upon expansion of the sum, note that any term in which an index appears only once will be zero due to the independence of the A_i 's and $\mathbb{E}A_i = 0$.

Hence,

$$\mathbb{E} \left| \sum_{i=1}^N A_i \right|^4 = \sum_{i=1}^N \mathbb{E}|A_i|^4 + 6 \sum_{i<j}^N \mathbb{E}|A_i|^2 \mathbb{E}|A_j|^2. \quad (4.84)$$

Letting M_4 denote the uniform bound on the fourth moments of the A_i 's and noting that $\max_{i=1,\dots,N} \mathbb{E}|A_i|^2 \leq M_4^{1/2}$, we arrive at

$$\mathbb{E} \left| \sum_{i=1}^N A_i \right|^4 \leq M_4 N + 3M_4^{1/2} N(N-1) \quad (4.85)$$

$$\leq (3M_4^{1/2} + M_4) N^2. \quad (4.86)$$

Scaling (4.86) by $1/N^4$ yields the result. \square

We will need the following conditional version of Lemma 4.6 in the next lemma. Let \mathcal{H} be a sub σ -algebra of \mathcal{F} . Let A_i , $i = 1, \dots, N$ be independent random variables conditioned on \mathcal{H} with $\mathbb{E}[A_i | \mathcal{H}] = 0$ a.s., and such that $\mathbb{E}[|A_i|^4 | \mathcal{H}]$ is almost surely bounded, uniformly in N . Then we can obtain the bound

$$\mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N A_i \right|^4 \mid \mathcal{H} \right] \leq \frac{C}{N^2} \quad \text{a.s.} \quad (4.87)$$

The technique to prove this is identical to Lemma 4.6. Let us now prove

the following auxiliary lemma.

Lemma 4.7. *Let Assumptions 4.3 and 4.5 be satisfied. Then the following bounds hold.*

For all $i = 1, \dots, N$,

$$\begin{aligned} & \mathbb{E}|b(t, X^{i,N}(t), \hat{\mu}_{\mathbf{X}^N(t)}) - b(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)})|^4 \\ & + \mathbb{E}|\sigma(t, X^{i,N}(t), \hat{\mu}_{\mathbf{X}^N(t)}) - \sigma(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)})|^4 \\ & \leq C \left(\mathbb{E}|X^{i,N}(t) - X^i(t)|^4 + \frac{1}{N} \sum_{j=1}^N \mathbb{E}|X^{j,N}(t) - X^j(t)|^4 \right), \end{aligned} \quad (4.88)$$

where $C > 0$ is independent of N .

For all $i = 1, \dots, N$,

$$\begin{aligned} & \mathbb{E}|b(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - b(t, X^i(t), \mathcal{L}_t^X)|^4 \\ & + \mathbb{E}|\sigma(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - \sigma(t, X^i(t), \mathcal{L}_t^X)|^4 \leq \frac{C}{N^2}, \end{aligned} \quad (4.89)$$

where $C > 0$ is independent of N .

Proof. Thanks to Assumption 4.5, we obtain

$$\begin{aligned} & \mathbb{E}|b(t, X^{i,N}(t), \hat{\mu}_{\mathbf{X}^N(t)}) - b(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)})|^4 \\ & + \mathbb{E}|\sigma(t, X^{i,N}(t), \hat{\mu}_{\mathbf{X}^N(t)}) - \sigma(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)})|^4 \\ & \leq C \mathbb{E}|X^{i,N}(t) - X^i(t)|^4 \\ & + C \sum_{j=1}^J \mathbb{E} \left| \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^{l,N}(t)) - \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^l(t)) \right|^4 \\ & \leq C \mathbb{E}|X^{i,N}(t) - X^i(t)|^4 + C \frac{1}{N} \sum_{l=1}^N \mathbb{E}|X^{l,N}(t) - X^l(t)|^4, \end{aligned}$$

where the final inequality follows from the fact that the ϕ_j 's are Lipschitz in their final argument. Hence, we have arrived at (4.88).

For the second bound, using Assumption 4.5 once again, we get

$$\begin{aligned}
& \mathbb{E}|b(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - b(t, X^i(t), \mathcal{L}_t^X)|^4 + \mathbb{E}|\sigma(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - \sigma(t, X^i(t), \mathcal{L}_t^X)|^4 \\
& \leq C \sum_{j=1}^J \mathbb{E} \left| \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^l(t)) - \int_G \phi_j(t, X^i(t), y) d\mathcal{L}_t^X(y) \right|^4 \\
& \leq C \sum_{j=1}^J \mathbb{E} \left| \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^l(t)) - \frac{1}{N-1} \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) \right|^4 \\
& \quad + C \sum_{j=1}^J \mathbb{E} \left| \frac{1}{N-1} \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) - \int_G \phi_j(t, X^i(t), y) \mathcal{L}_t^X(dy) \right|^4.
\end{aligned} \tag{4.90}$$

Now, for the first term in (4.90), note that

$$\begin{aligned}
& \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^l(t)) - \frac{1}{N-1} \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) \\
& = \left(\frac{1}{N} - \frac{1}{N-1} \right) \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) + \frac{1}{N} \phi_j(t, X^i(t), X^i(t)).
\end{aligned} \tag{4.91}$$

So, under Assumption 4.3, we have

$$\begin{aligned}
& \mathbb{E} \left| \frac{1}{N} \sum_{l=1}^N \phi_j(t, X^i(t), X^l(t)) - \frac{1}{N-1} \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) \right|^4 \\
& \leq \frac{C_1}{N^8} + \frac{C_2}{N^4}.
\end{aligned}$$

Hence, we have

$$\begin{aligned}
& \mathbb{E}|b(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - b(t, X^i(t), \mathcal{L}_t^X)|^4 + \mathbb{E}|\sigma(t, X^i(t), \hat{\mu}_{\mathbf{X}(t)}) - \sigma(t, X^i(t), \mathcal{L}_t^X)|^4 \\
& \leq \frac{C}{N^4} + C \sum_{j=1}^J \mathbb{E} \left(\mathbb{E} \left(\left| \frac{1}{N-1} \sum_{l=1, l \neq i}^N \phi_j(t, X^i(t), X^l(t)) \right. \right. \right. \\
& \quad \left. \left. \left. - \int_G \phi_j(t, X^i(t), y) \mathcal{L}_t^X(dy) \right|^4 \middle| X^i(t) \right) \right).
\end{aligned} \tag{4.92}$$

Finally, note that for each j

$$A_j^l := \phi_j(t, X^i(t), X^l(t)) - \int_G \phi_j(t, X^i(t), y) \mathcal{L}_t^X(\mathbf{d}y), \quad (4.93)$$

for $l = 1, \dots, N$, $l \neq i$ are conditionally independent random variables, conditioned on $X^i(t)$, with uniformly bounded moments and $\mathbb{E}[A_j^l | X^i(t)] = 0$ a.s.. Then, applying the conditional version (4.87) of Lemma 4.6, we arrive at (4.89). \square

Now we proceed to the propagation of chaos theorem.

Theorem 4.8. *Let Assumptions 4.3 and 4.5 be satisfied. Then the following bound holds:*

$$\left(\max_{i=1, \dots, N} \mathbb{E} \sup_{u \leq t} |X^{i,N}(u) - X^i(u)|^4 \right)^{1/4} \leq \frac{C}{N^{1/2}}, \quad (4.94)$$

where $X^{i,N}$ represents the i -th particle among the interacting particles driven by (4.19), X^i is independent and identical copy of (4.17), and $C > 0$ is a constant independent of N .

Proof. For brevity, we write $X_t^{i,N} := X^{i,N}(t)$ and $X_t^i := X^i(t)$. Analogously to how (4.55) was derived, we can get

$$\begin{aligned} \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4 &\leq K \mathbb{E} \left(\int_0^t |X_s^{i,N} - X_s^i|^2 |b(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - b(s, X_s^i, \mathcal{L}_s^X)|^2 \right. \\ &\quad + |X_s^{i,N} - X_s^i|^2 |\sigma(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - \sigma(s, X_s^i, \mathcal{L}_s^X)|^2 \\ &\quad \left. + |X_s^{i,N} - X_s^i|^4 + |\sigma(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - \sigma(s, X_s^i, \mathcal{L}_s^X)|^4 \mathbf{d}s \right). \end{aligned} \quad (4.95)$$

Using Young's inequality, we have

$$\begin{aligned} \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4 &\leq K \mathbb{E} \left(\int_0^t |X_s^{i,N} - X_s^i|^4 + |b(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - b(s, X_s^i, \mathcal{L}_s^X)|^4 \right. \\ &\quad \left. + |\sigma(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - \sigma(s, X_s^i, \mathcal{L}_s^X)|^4 \mathbf{d}s \right). \end{aligned} \quad (4.96)$$

We also have

$$|b(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - b(s, X_s^i, \mathcal{L}_s^X)|^4 \leq 8(|b(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - b(s, X_s^i, \hat{\mu}_{\mathbf{X}_s})|^4 + |b(s, X_s^i, \hat{\mu}_{\mathbf{X}_s}) - b(s, X_s^i, \mathcal{L}_s^X)|^4), \quad (4.97)$$

$$|\sigma(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - \sigma(s, X_s^i, \mathcal{L}_s^X)|^4 \leq 8(|\sigma(s, X_s^{i,N}, \hat{\mu}_{\mathbf{X}_s^N}) - \sigma(s, X_s^i, \hat{\mu}_{\mathbf{X}_s})|^4 + |\sigma(s, X_s^i, \hat{\mu}_{\mathbf{X}_s}) - \sigma(s, X_s^i, \mathcal{L}_s^X)|^4). \quad (4.98)$$

Using (4.97), (4.98) and Lemma 4.7, we obtain from (4.96):

$$\begin{aligned} \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4 &\leq C \int_0^t \mathbb{E} |X^{i,N}(s) - X^i(s)|^4 \\ &\quad + \frac{1}{N} \sum_{j=1}^N \mathbb{E} |X^{j,N}(s) - X^j(s)|^4 ds + C \frac{1}{N^2}. \end{aligned}$$

Now, using the notation

$$\begin{aligned} \Delta_t^i &:= \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4, \\ \Delta_t &:= \max_{i=1, \dots, N} \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4, \end{aligned}$$

we have from (4.99)

$$\begin{aligned} \Delta_t^i &\leq C \int_0^t \left(\Delta_s^i + \frac{1}{N} \sum_{j=1}^N \Delta_s^j \right) ds + C \frac{1}{N^2} \\ &\leq C \int_0^t \Delta_s ds + C \frac{1}{N^2}. \end{aligned}$$

From which we can take the maximum over i to obtain

$$\Delta_t \leq C \int_0^t \Delta_s ds + C \frac{1}{N^2}. \quad (4.99)$$

Application of Grönwall's lemma gives us the desired result

$$\Delta_t = \max_{i=1, \dots, N} \mathbb{E} \sup_{u \leq t} |X_u^{i,N} - X_u^i|^4 \leq \frac{C}{N^2}. \quad (4.100)$$

□

We will verify that the CBO models from Section 4.3.1 satisfy Assumption 4.5. Consequently, we get the following corollary.

Corollary 4.9. *The particle systems (4.26) and (4.29) converge to their mean-field limits (4.24) and (4.28), respectively: for them the inequality (4.94) holds.*

To see why the CBO models from Section 4.3.1 satisfy Assumption 4.5, consider the following argument, which follows the technique from [28]. We use the following notation: for a probability measure $\mu \in \mathcal{P}(\bar{G})$, define

$$A(\mu) := \int_{\bar{G}} x e^{-\alpha f(x)} \mu(\mathrm{d}x), \quad (4.101)$$

$$B(\mu) := \int_{\bar{G}} e^{-\alpha f(x)} \mu(\mathrm{d}x), \quad (4.102)$$

so that $\bar{X}^\mu = A(\mu)/B(\mu)$. Then, for two measures $\mu_1, \mu_2 \in \mathcal{P}(\bar{G})$

$$\begin{aligned} |\bar{X}^{\mu_1} - \bar{X}^{\mu_2}| &= \left| \frac{A(\mu_1)}{B(\mu_1)} - \frac{A(\mu_2)}{B(\mu_2)} \right| \\ &= \frac{|A(\mu_1)B(\mu_2) - A(\mu_2)B(\mu_1)|}{B(\mu_1)B(\mu_2)}. \end{aligned} \quad (4.103)$$

Letting $m := \inf_{x \in \bar{G}} e^{-\alpha f(x)}$, we have

$$\begin{aligned} |\bar{X}^{\mu_1} - \bar{X}^{\mu_2}| &\leq \frac{|A(\mu_1)||B(\mu_2) - B(\mu_1)| + |B(\mu_1)||A(\mu_1) - A(\mu_2)|}{m^2}, \\ &\leq C(|B(\mu_1) - B(\mu_2)| + |A(\mu_1) - A(\mu_2)|), \end{aligned} \quad (4.104)$$

for some $C > 0$, where we use the fact that G is bounded. For

$$\begin{aligned} b(t, x, \mu) &= -\beta(x - \bar{X}^\mu), \\ \sigma(t, x, \mu) &= \sigma \operatorname{Diag}(x - \bar{X}^\mu), \end{aligned}$$

we have

$$\begin{aligned} &|b(t, x, \mu_1) - b(t, y, \mu_2)| + |\sigma(t, x, \mu_1) - \sigma(t, y, \mu_2)| \\ &= (\beta + \sigma)|x - y| + (\bar{X}^{\mu_2} - \bar{X}^{\mu_1})| \\ &\leq (\beta + \sigma)(|x - y| + |\bar{X}^{\mu_1} - \bar{X}^{\mu_2}|), \end{aligned} \quad (4.105)$$

which when combined with the bound (4.104) shows that the coefficients satisfy Assumption 4.5 with $J = 2$ and

$$\begin{aligned}\phi_1(t, x, z) &= e^{-\alpha f(z)}, \\ \phi_2(t, x, z) &= ze^{-\alpha f(z)}.\end{aligned}$$

For the repelling system, (4.28), we follow the same argument but need an additional function

$$\phi_3(t, x, z) = \phi_3(x, z) = (x - z)e^{-\frac{1}{2}|x-z|^2},$$

in order to deal with the repelling term in the drift coefficient. In (4.28), we have

$$\begin{aligned}b(t, x, \mu) &= -\beta(x - \bar{X}^\mu) + \lambda(t)H(x, \mu), \\ \sigma(t, x, \mu) &= \sigma \text{Diag}(x - \bar{X}^\mu),\end{aligned}$$

where

$$H(x, \mu) := \int_G (x - z)e^{-\frac{1}{2}|x-z|^2} \mu(\mathbf{d}z). \quad (4.106)$$

As such, we have

$$\begin{aligned}& |b(t, x, \mu_1) - b(t, y, \mu_2)| + |\sigma(t, x, \mu_1) - \sigma(t, y, \mu_2)| \\ & \leq (\beta + \sigma)(|x - y| + |\bar{X}^{\mu_1} - \bar{X}^{\mu_2}|) + \lambda(0)|H(x, \mu_1) - H(y, \mu_2)|,\end{aligned} \quad (4.107)$$

recalling that $\lambda(t)$ is non-negative and decreasing. To deal with the final term in (4.107), notice that

$$\begin{aligned}|H(x, \mu_1) - H(y, \mu_2)| &\leq \left| \int_G (\phi_3(x, z) - \phi_3(y, z)) \mu_1(\mathbf{d}z) \right| \\ & \quad + \left| \int_G \phi_3(y, z) \mu_1(\mathbf{d}z) - \int_G \phi_3(y, z) \mu_2(\mathbf{d}z) \right|.\end{aligned} \quad (4.108)$$

The second term in (4.108) is precisely the bound we need in (4.79) of

Assumption 4.5. Meanwhile,

$$\begin{aligned} \left| \int_G (\phi_3(x, z) - \phi_3(y, z)) \mu_1(\mathbf{d}z) \right| &\leq \left| \int_G C|x - y| \mu_1(\mathbf{d}z) \right| \\ &= C|x - y|, \end{aligned} \quad (4.109)$$

due to ϕ_3 being Lipschitz in each argument (we will show this momentarily). Thus,

$$\begin{aligned} |H(x, \mu_1) - H(y, \mu_2)| &\leq C|x - y| \\ &\quad + \left| \int_G \phi_3(y, z) (\mu_1 - \mu_2)(\mathbf{d}z) \right|, \end{aligned} \quad (4.110)$$

which when combined with (4.107) shows that the repelling system's coefficients satisfy Assumption 4.5.

To verify that ϕ_3 is Lipschitz, fix $z \in \bar{G}$ and consider the Jacobian of ϕ_3 with respect to x , which we denote by J_{ϕ_3} . We have

$$J_{\phi_3}(x) = (I_d - (x - z)(x - z)^\top) e^{-\frac{1}{2}|x - z|^2}. \quad (4.111)$$

Letting $\|\cdot\|_2$ denote the operator norm induced by the Euclidean norm, we have

$$\begin{aligned} \|J_{\phi_3}\|_2 &\leq e^{-\frac{1}{2}|x - z|^2} (\|I_d\|_2 + \|(x - z)(x - z)^\top\|_2) \\ &\leq 1 + |x - z|^2 \\ &\leq 1 + 4M^2, \end{aligned} \quad (4.112)$$

where $M := \sup_{x \in \bar{G}} |x|$, recalling Assumption 4.3 that G is bounded. Now,

$$\begin{aligned} \phi_3(x, z) - \phi_3(y, z) &= \int_0^1 J_{\phi_3}(y + s(x - y))(x - y) \mathbf{d}s \\ &\leq \int_0^1 (1 + 4M^2) |x - y| \mathbf{d}s \\ &= (1 + 4M^2) |x - y|. \end{aligned}$$

Note that the path $s \rightarrow y + s(x - y)$ may not be contained in \bar{G} since we do

not assume that G is convex; however, we can take x, y from the convex hull of \bar{G} and consider the extension of ϕ_3 to this space: the analysis remains the same and ϕ_3 is Lipschitz in the first argument.

4.6 Convergence of mean-field limit of CBO-type models to the global minimum

This section is devoted to showing convergence of the constrained CBO models from Section 4.3.1 towards global minimizers. Similar to [28, 30] (see also [71]), the main tool which we will utilize is the Laplace principle which states that for any compactly supported measure μ and positive function f the following holds:

$$\lim_{\alpha \rightarrow \infty} -\frac{1}{\alpha} \log \int_{\mathbb{R}^d} e^{-\alpha f(x)} \mu(\mathrm{d}x) = f(x_{\min}) \quad (4.113)$$

with x_{\min} being the minimizer of $f(x)$ defined on the support of μ . We will use the notation

$$f_{\text{osc}} = f_{\max} - f_{\min}.$$

The analysis in this section is done for convex domains.

Assumption 4.6. *G is a convex bounded domain.*

Let us explain why we restrict analysis in this section to convex domains. In the case of non-convex domains, CBO models with constraints of the type presented in Section 4.3.1 face two difficulties. The first difficulty concerns the proof techniques used for showing convergence to the global minimum which have been used so far in the CBO literature (see, e.g. [28, 30, 71, 8] and references therein and also the proofs below in this section). Roughly speaking, noise does not play a positive role within these proofs which rely on the drift driving solutions of the mean-field SDEs to the consensus. It works well in the convex case and in the whole Euclidean space. However, in situations like the one illustrated in Figure 2, the drift forces trajectories to go outside the domain rather than driving it to a consensus and, hence, we need to rely on the noise term

(via effectively large deviations) to get trajectories to a proximity of the global minimum, which happens with a positive probability, and then the drift can drive trajectories to a consensus. Hence, a probabilistic proof needs to be developed for proving convergence to the global minimum, which is an interesting problem for future research. Such a proof would also be able to provide more insight on how convergence to the global minimum depends on the noise strength, which would be useful even in the whole Euclidean space. The second difficulty is more fundamental, as it is related to the design of CBO models themselves. In the example of Figure 2, the drift forces trajectories the wrong way because its form is based on the Euclidean metric. Therefore, it is of interest to find a different formulation of CBO models which imbibe more information of the domain (e.g., by using a non-Euclidean metric) but at the same time remain computationally efficient. Constructing such models is out of scope of this work. Nevertheless, we tested the CBO model given by Equation (4.26) on a non-convex function with non-convex domain (see Section 4.7.3) and the CBO model demonstrated excellent performance.

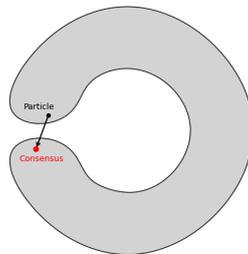


Figure 2: The example of a non-convex domain, where the Euclidean path of a particle to the consensus does not lie within the domain. In the convex case, this behaviour cannot occur.

In Theorem 4.10, we prove convergence of the reflected CBO model, Equation (4.24), towards the global minimum and consider the case of CBO with repelling force, Equation (4.28), in Theorem 4.11. We note that under Assumptions 4.4 and 4.6 well-posedness of (4.17) and the corresponding propagation of chaos were proved in [1].

Theorem 4.10. *Let Assumptions 4.1 and 4.6 hold. If the following condition is satisfied*

$$\eta_0 := 2\beta - \sigma^2(1 + e^{2\alpha f_{osc}}) > 0 \quad (4.114)$$

then there exists an $x^ \in \bar{G}$ such that $X(t) \rightarrow x^*$ almost surely as $t \rightarrow \infty$ and*

$$f(x^*) \leq f_{\min} + \Gamma(\alpha),$$

where $X(t)$ is the solution of the mean-field SDE (4.24) and $\Gamma(\alpha) \rightarrow 0$ as $\alpha \rightarrow \infty$.

Remark 4.1. *Theorem 4.10 has two parts. First, the solution of the mean-field SDE (4.24) converges to some x^* almost surely as time goes to infinity. Second, the closeness of this limit point x^* to the minimizer is controlled by α , in the sense that $f(x^*)$ can be made arbitrarily close to f_{\min} provided α is large enough. Here, α has the interpretation of the inverse temperature and governs how selective the consensus point is: the consensus point becomes dominated by the best-performing particle(s) as $\alpha \rightarrow \infty$.*

Before we prove Theorem 4.10, let us remark about the convexity of G . First, since we are no longer assuming smoothness of ∂G , it is possible that we have points on ∂G where the normal direction is not unique. For example, the corners of a convex polytope. From now on, $\nu(x), x \in \partial G$ refers to a choice of inward normal (if there is more than one). Additionally, we will use the following property of convex domains, which can be found in e.g. [1]. For any $x \in \partial G$ and any $y \in \bar{G}$, we have

$$\nu(x) \cdot (y - x) \geq 0. \quad (4.115)$$

We make use of this in the proof of Theorems 4.10 and 4.11.

Proof. Using Ito's formula, we get

$$\begin{aligned}
|X(t) - \mathbb{E}X(t)|^2 &= |X(0) - \mathbb{E}X(0)|^2 \\
&\quad - 2 \int_0^t \beta(X(s) - \mathbb{E}X(s)) \cdot (X(s) - \bar{X}(s)) \mathbf{d}s \\
&\quad - 2 \int_0^t (X(s) - \mathbb{E}X(s))^\top \mathbf{d}\mathbb{E}X(s) \\
&\quad + 2 \int_0^t \sigma((X(s) - \mathbb{E}X(s)) \cdot \text{Diag}(X(s) - \bar{X}(s))) \mathbf{d}W(s) \\
&\quad + \int_0^t \sigma^2 |X(s) - \bar{X}(s)|^2 \mathbf{d}s + 2 \int_0^t ((X(s) - \mathbb{E}X(s)) \cdot \nu(X(s))) \mathbf{d}L(s) \\
&= |X(0) - \mathbb{E}X(0)|^2 - 2 \int_0^t \beta |X(s) - \mathbb{E}X(s)|^2 \mathbf{d}s \\
&\quad - 2 \int_0^t (X(s) - \mathbb{E}X(s))^\top \mathbf{d}\mathbb{E}X(s) \\
&\quad - 2 \int_0^t \beta((X(s) - \mathbb{E}X(s))) \cdot (\mathbb{E}X(s) - \bar{X}(s)) \mathbf{d}s \\
&\quad + \int_0^t \sigma^2 |X(s) - \bar{X}(s)|^2 \mathbf{d}s \\
&\quad + 2 \int_0^t \sigma((X(s) - \mathbb{E}X(s)) \cdot \text{Diag}(X(s) - \bar{X}(s))) \mathbf{d}W(s) \\
&\quad + 2 \int_0^t ((X(s) - \mathbb{E}X(s)) \cdot \nu(X(s))) \mathbf{d}L(s).
\end{aligned}$$

Due to Assumption 4.6, we have

$$(X(s) - \mathbb{E}X(s)) \cdot \nu(X(s)) \leq 0.$$

Also, note that $\mathbb{E}((X(s) - \mathbb{E}X(s)) \cdot (\mathbb{E}X(s) - \bar{X}(s))) = 0$. Consequently, we get

$$\begin{aligned}
\mathbb{V}(t) := \mathbb{E}|X(t) - \mathbb{E}X(t)|^2 &\leq \mathbb{E}|X_0 - \mathbb{E}(X_0)|^2 - 2 \int_0^t \beta \mathbb{E}|X(s) - \mathbb{E}X(s)|^2 \mathbf{d}s \\
&\quad + \int_0^t \sigma^2 \mathbb{E}|X(s) - \bar{X}(s)|^2 \mathbf{d}s.
\end{aligned}$$

We have

$$\mathbb{E}|X(s) - \bar{X}(s)|^2 = \mathbb{E}|X(s) - \mathbb{E}X(s)|^2 + |\mathbb{E}X(s) - \bar{X}(s)|^2 \quad (4.116)$$

and

$$\begin{aligned}
|\mathbb{E}X(s) - \bar{X}(s)|^2 &= \left| \mathbb{E}X(s) - \frac{\mathbb{E}X(s)e^{-\alpha f(X(s))}}{\mathbb{E}e^{-\alpha f(X(s))}} \right|^2 \\
&= \left| \mathbb{E} \left(\left(\mathbb{E}X(s) - X(s) \right) \frac{e^{-\alpha f(X(s))}}{\mathbb{E}e^{-\alpha f(X(s))}} \right) \right|^2 \\
&\leq e^{2\alpha f_{\text{osc}}} \mathbb{E}|X(s) - \mathbb{E}X(s)|^2.
\end{aligned} \tag{4.117}$$

Therefore,

$$\frac{d}{dt} \mathbb{V}(t) \leq \left(-2\beta + \sigma^2(1 + e^{2\alpha f_{\text{osc}}}) \right) \mathbb{V}(t), \tag{4.118}$$

which, using (4.114), implies

$$\mathbb{V}(t) \leq \mathbb{V}(0)e^{-\eta_0 t}. \tag{4.119}$$

This also means, due to (4.116) and (4.117), there exist positive constants c_1 and c_2 independent of t such that

$$\mathbb{E}|X(t) - \bar{X}(t)|^2 \leq c_1 e^{-\eta_0 t}, \tag{4.120}$$

$$|\mathbb{E}X(t) - \bar{X}(t)|^2 \leq c_2 e^{-\eta_0 t}. \tag{4.121}$$

Let us fix a constant $q \in G$. Then, again using Ito's formula, we have

$$\begin{aligned}
d|X(t) - q|^2 &= -2\beta((X(t) - q) \cdot (X(t) - \bar{X}(t)))dt + \sigma^2|X(t) - \bar{X}(t)|^2 dt \\
&\quad + 2\sigma((X(t) - q) \cdot \text{Diag}(X(t) - \bar{X}(t))dW(t)) \\
&\quad + 2((X(t) - q) \cdot \nu(X(t)))dL(t).
\end{aligned}$$

Using Tanaka's trick (thanks to Assumption 4.6) and taking expectation on both sides, we have

$$\begin{aligned}
d\mathbb{E}|X(t) - q|^2 &\leq -2\beta\mathbb{E}((X(t) - q) \cdot (X(t) - \bar{X}(t)))dt + \sigma^2\mathbb{E}|X(t) - \bar{X}(t)|^2 dt \\
&\leq -2\beta\mathbb{E}((X(t) - q) \cdot (X(t) - \bar{X}(t)))dt + \sigma^2 c_1 e^{-\eta_0 t} dt,
\end{aligned}$$

where we have used (4.120) in the last inequality. Splitting the term

$((X(t) - q) \cdot (X(t) - \bar{X}(t)))$, we get

$$\begin{aligned} \frac{d}{dt} \mathbb{E}|X(t) - q|^2 &\leq -2\beta \mathbb{E}((X(t) - q) \cdot (X(t) - \bar{X}(t))) + \sigma^2 c_1 e^{-\eta_0 t} \\ &= -2\beta \mathbb{E}|X(t) - q|^2 + 2\beta((\mathbb{E}X(t) - q) \cdot (\bar{X}(t) - q)) \\ &\quad + \sigma^2 c_1 e^{-\eta_0 t}. \end{aligned}$$

Rewriting $((\mathbb{E}X(t) - q) \cdot (\bar{X}(t) - q)) = |\mathbb{E}X(t) - q|^2 + ((\mathbb{E}X(t) - q) \cdot (\bar{X}(t) - \mathbb{E}X(t)))$ gives

$$\begin{aligned} \frac{d}{dt} \mathbb{E}|X(t) - q|^2 &\leq -2\beta \mathbb{E}|X(t) - q|^2 + 2\beta |\mathbb{E}X(t) - q|^2 \\ &\quad + 2\beta((\mathbb{E}X(t) - q) \cdot (\bar{X}(t) - \mathbb{E}X(t))) + \sigma^2 c_1 e^{-\eta_0 t}. \end{aligned}$$

Due to Jensen's inequality, we get $|\mathbb{E}X(t) - q|^2 \leq \mathbb{E}|X(t) - q|^2$ which provides the following differential inequality:

$$\begin{aligned} \frac{d}{dt} \mathbb{E}|X(t) - q|^2 &\leq 2\beta((\mathbb{E}X(t) - q) \cdot (\bar{X}(t) - \mathbb{E}X(t))) + \sigma^2 c_1 e^{-\eta_0 t} \\ &\leq 2\beta \sqrt{c_2} |\mathbb{E}X(t) - q| e^{-\eta_0 t/2} + \sigma^2 c_1 e^{-\eta_0 t}. \end{aligned}$$

Note that $|\mathbb{E}X(t) - q|$ is bounded uniformly in t due to boundedness of the domain G . Therefore, we can say that there exists a constant $c_3 > 0$ such that

$$\frac{d}{dt} \mathbb{E}|X(t) - q|^2 \leq c_3 e^{-\eta_0 t/2}, \quad (4.122)$$

which implies that there exists an $x^* \in \bar{G}$ such that $\mathbb{E}X(t)$ is converging to x^* as $t \rightarrow \infty$. Here, x^* belongs to \bar{G} due to convexity of G .

Using Markov's inequality, we get

$$\mathbb{P}(|X(t) - \mathbb{E}X(t)| \geq e^{-\eta_0 t/4}) \leq \frac{\mathbb{V}(t)}{e^{-\eta_0 t/2}} \leq C e^{-\eta_0 t/2},$$

where C is a positive constant independent of t . Then, using the Borel-Cantelli lemma, $|X(t) - \mathbb{E}X(t)| \rightarrow 0$ as $t \rightarrow \infty$ a.s., and hence $X(t) \rightarrow x^*$ a.s.. Consequently, applying the bounded convergence theorem, we

arrive at $\mathbb{E}e^{-\alpha f(X(t))} \rightarrow e^{-\alpha f(x^*)}$ as $t \rightarrow \infty$. Therefore, we can say: for all $\epsilon > 0$ there exists a (deterministic) $T^* > 0$ such that

$$|\mathbb{E}e^{-\alpha f(X(t))} - e^{-\alpha f(x^*)}| \leq \epsilon \quad (4.123)$$

for all $t \geq T^*$. Since the probability distribution of $X(T^*)$ is compactly supported, using the Laplace principle we obtain

$$-\frac{1}{\alpha} \log(\mathbb{E}e^{-\alpha f(X(T^*))}) \leq f_{\min} + \Gamma_1(\alpha), \quad (4.124)$$

where $\Gamma_1(\alpha) \rightarrow 0$ as $\alpha \rightarrow \infty$. Note that the exact form of Γ_1 is neither specified nor important. The statement (4.124) simply expresses that the left-hand side of (4.124) decreases to f_{\min} , this is exactly the Laplace principle. And, from (4.123), we have

$$\mathbb{E}e^{-\alpha f(X(T^*))} \leq \epsilon + e^{-\alpha f(x^*)},$$

which implies

$$\begin{aligned} \log \mathbb{E}e^{-\alpha f(X(T^*))} &\leq \log(\epsilon + e^{-\alpha f(x^*)}), \\ -\frac{1}{\alpha} \log \mathbb{E}e^{-\alpha f(X(T^*))} &\geq -\frac{1}{\alpha} \log(\epsilon + e^{-\alpha f(x^*)}). \end{aligned} \quad (4.125)$$

Using (4.124) and (4.125), we obtain

$$-\frac{1}{\alpha} \log(\epsilon + e^{-\alpha f(x^*)}) \leq f_{\min} + \Gamma_1(\alpha). \quad (4.126)$$

From the mean value theorem, we have

$$\log(\epsilon + e^{-\alpha f(x^*)}) = -\alpha f(x^*) + \frac{\epsilon}{\gamma\epsilon + e^{-\alpha f(x^*)}}, \quad (4.127)$$

for some $\gamma \in (0, 1)$. Therefore,

$$f(x^*) \leq f_{\min} + \Gamma_1(\alpha) + \frac{\epsilon}{\alpha \gamma \epsilon + e^{-\alpha f(x^*)}}.$$

□

We now consider the repelling CBO model (4.28) and prove an analogous result.

Theorem 4.11. *Let Assumptions 4.1 and 4.6 hold. Let $X(t)$ be the solution to (4.28) and let*

$$\eta_1 := -2\beta + 3\lambda(0) + \sigma^2(1 + e^{2\alpha f_{osc}}) \quad (4.128)$$

be strictly greater than a positive constant K_1 . If $0 \leq \lambda(t) \leq K_2 \exp(-\eta_2 t)$ for some $K_2 > 0$ and $\eta_2 > \eta_1$, then there exists an $x^ \in \bar{G}$ such that $X(t) \rightarrow x^*$ as $t \rightarrow \infty$ and*

$$f(x^*) \leq f_{\min} + \Gamma(\alpha),$$

where $\Gamma(\alpha) \rightarrow 0$ as $\alpha \rightarrow \infty$.

Proof. Using Ito's formula, we have

$$\begin{aligned} |X(t) - \mathbb{E}X(t)|^2 &= |X(0) - \mathbb{E}X(0)|^2 \\ &\quad - 2 \int_0^t \beta(X(s) - \mathbb{E}X(s)) \cdot (X(s) - \bar{X}(s)) \mathbf{d}s \\ &\quad + 2 \int_0^t \lambda(s) \int_{\mathbb{R}^d} ((X(s) - \mathbb{E}X(s)) \cdot (X(s) - y)) \\ &\quad \times \exp\left(-\frac{1}{2}|X(s) - y|^2\right) \mathcal{L}_{X(s)}(\mathbf{d}y) \mathbf{d}s \\ &\quad + 2 \int_0^t \sigma((X(s) - \mathbb{E}X(s)) \cdot \text{Diag}(X(s) - \bar{X}(s))) \mathbf{d}W(s) \\ &\quad + \int_0^t \sigma^2 |X(s) - \bar{X}(s)|^2 \mathbf{d}s + 2 \int_0^t ((X(s) - \mathbb{E}X(s)) \cdot \nu(X(s))) \mathbf{d}L(s). \end{aligned}$$

The objective here is to deal with the repelling term. In that pursuit, we

apply Young's inequality to get

$$\begin{aligned}
\mathbb{E} \int_{\mathbb{R}^d} ((X(s) - \mathbb{E}X(s)) \cdot (X(s) - y)) \exp\left(-\frac{1}{2}|X(s) - y|^2\right) \mathcal{L}_{X(s)}(\mathbf{d}y) \\
\leq \frac{1}{2} \mathbb{E}|X(s) - \mathbb{E}X(s)|^2 \\
+ \frac{1}{2} \mathbb{E} \int_{\mathbb{R}^d} |X(s) - y|^2 \exp\left(-\frac{1}{2}|X(s) - y|^2\right) \mathcal{L}_{X(s)}(\mathbf{d}y) \\
\leq \frac{1}{2} \mathbb{E}|X(s) - \mathbb{E}X(s)|^2 + \frac{1}{2} \mathbb{E} \int_{\mathbb{R}^d} |X(s) - y|^2 \mathcal{L}_{X(s)}(\mathbf{d}y) \\
\leq \frac{3}{2} \mathbb{E}|X(s) - \mathbb{E}X(s)|^2, \tag{4.129}
\end{aligned}$$

since $\mathbb{E} \int_{\mathbb{R}^d} |X(s) - y|^2 \mathcal{L}_{X(s)}(\mathbf{d}y) = 2\mathbb{E}|X(s) - \mathbb{E}X(s)|^2$.

Exploiting similar arguments as the ones used to get (4.118) and using the inequality obtained in (4.129), we ascertain

$$\frac{\mathbf{d}}{\mathbf{d}t} \mathbb{V}(t) \leq \left(-2\beta + 3\lambda(0) + \sigma^2(1 + e^{2\alpha f_{\text{osc}}})\right) \mathbb{V}(t).$$

With β , σ and $\lambda(0)$ satisfying (4.128), we have

$$\frac{\mathbf{d}}{\mathbf{d}t} \mathbb{V}(t) \leq c_4 e^{-\eta_1 t},$$

where c_4 and η_1 are independent of t .

Let us again fix a constant $q \in G$. Applying Ito's formula gives

$$\begin{aligned}
\mathbf{d}|X(t) - q|^2 &= -2\beta((X(t) - q) \cdot (X(t) - \bar{X}(t))) \mathbf{d}t + \sigma^2 |X(t) - \bar{X}(t)|^2 \mathbf{d}t \\
&+ 2\lambda(t) \int_{\mathbb{R}^d} ((X(t) - q) \cdot (X(t) - y)) \\
&\quad \times \exp\left(-\frac{1}{2}|X(t) - y|^2\right) \mathcal{L}_{X(t)}(\mathbf{d}y) \mathbf{d}t \\
&+ 2\sigma((X(t) - q) \cdot \text{Diag}(X(t) - \bar{X}(t)) \mathbf{d}W(t)) \\
&+ 2(X(t) - q) \cdot \nu(X(t)) \mathbf{d}L(t).
\end{aligned}$$

Using Young's inequality, we have

$$\begin{aligned}
& \int_{\mathbb{R}^d} ((X(t) - q) \cdot (X(t) - y)) \exp\left(-\frac{1}{2}|X(t) - y|^2\right) \mathcal{L}_{X(t)}(\mathbf{d}y) \\
&= \frac{1}{2}|X(t) - q|^2 \int_{\mathbb{R}^d} \exp\left(-\frac{1}{2}|X(t) - y|^2\right) \mathcal{L}_{X(t)}(\mathbf{d}y) \\
&\quad + \frac{1}{2} \int_{\mathbb{R}^d} |X(t) - y|^2 \exp\left(-\frac{1}{2}|X(t) - y|^2\right) \mathcal{L}_{X(t)}(\mathbf{d}y) \\
&\leq \frac{1}{2}|X(t) - q|^2 + \int_{\mathbb{R}^d} \frac{1}{2}|X(t) - y|^2 \mathcal{L}_{X(t)}(\mathbf{d}y).
\end{aligned}$$

Using analogous arguments employed to obtain (4.122), we get

$$\begin{aligned}
\frac{\mathbf{d}}{\mathbf{d}t} \mathbb{E}|X(t) - q|^2 &\leq c_5 e^{-\eta t/2} \lambda(t) \mathbb{E}|X(t) - q|^2 \\
&\quad + \lambda(t) \mathbb{E} \int_{\mathbb{R}^d} |X(t) - y|^2 \mathcal{L}_{X(t)}(\mathbf{d}y) \\
&\leq c_5 e^{-\eta t/2} + \lambda(t) \mathbb{E}|X(t) - q|^2 + 2\lambda(t)c_4 e^{-\eta t}, \quad (4.130)
\end{aligned}$$

where $c_5 > 0$ is a constant independent of t . Note that $\mathbb{E}|X(t) - q|^2$ is bounded uniformly in t . Since the function $\lambda(t)$ tends to zero faster than $e^{-\eta t}$ as $t \rightarrow \infty$, the right-hand side of (4.130) is integrable and we get that $\mathbb{E}|X(t) - q|^2$ converges to a constant as $t \rightarrow \infty$. This implies that $|\mathbb{E}X(t) - q|^2$ converges to a constant for every g in G . As in the proof of Theorem 4.10, this implies that there is an $x^* \in \bar{G}$ such that $\mathbb{E}X(t) \rightarrow x^*$ as $t \rightarrow \infty$. The rest of the arguments to complete the proof are the same as used in the previous theorem. \square

4.7 Numerical tests

In this section, we first (Section 4.7.1) consider discrete approximations of the interacting particle systems and then perform several numerical experiments to demonstrate effectiveness of the CBO models described in Section 4.3.1 for constrained optimization.

Since these CBO algorithms are probabilistic in nature and convergence is guaranteed only asymptotically (in N , α and t), we report the success rate of each experiment for a finite N , α and t . Each experiment is run 10^3 times (except the experiment in Section 4.7.3, which is only run

100 times), and the success rate is simply the proportion of experiments which successfully find the true global minimizer. We declare an experiment successful if the final consensus is within $\varepsilon > 0$ of the location of the true minimizer. For each experiment, we use the threshold $\varepsilon = 0.1$ unless otherwise specified.

We note briefly that there are several parameters/controls for the CBO algorithms: (i) the inverse temperature, α ; (ii) the number of particles, N ; (iii) strength of the drift and diffusion coefficients, β and σ , respectively; (iv) the integration time, t ; and (v) the step size, h , of a numerical scheme approximating the particle system. We do not attempt to propose a systematic way of choosing these parameters, but instead opt for sensible and practically effective choices in the following experiments. In general, α and N should be chosen as large as possible. For α , we run the risk of numerical instability if it is chosen too large. For N , the computational cost of the algorithm is at least $\mathcal{O}(N)$ (and $\mathcal{O}(N^2)$ in the case of the repelling CBO defined by Equation (4.24)), so there is a trade-off that must be balanced. As such, in Sections 4.7.2, 4.7.3, 4.7.4, and 4.7.5 we explore success rates for varying values of N . Similar comments can be made about the number of steps of the numerical scheme.

Note that in Theorems 4.10 and 4.11, the idea is that for a fixed α there exists some deterministic $T > 0$ such that the solution of the MVSDE is ε -close to its limit. Then the closeness of this limit to the optimum is controlled by α . So the time T that we should run the system for should depend on α . Here we do not propose a method to determine how to choose T or propose a method to determine a stopping time. Although the development of a stopping rule would be useful.

The strength of drift, β , and diffusion σ , coefficients can also depend on time. We experiment with both of them being constant in Sections 4.7.2, 4.7.3, and 4.7.5 but also linearly increasing (β) and exponentially decreasing (σ) functions in Sections 4.7.4 and 4.7.6. In all the experiments for the penalty scheme (4.131) given below the optimal choice of the penalty $\epsilon = h$ was used. Where possible, we choose the feasible region such

that the optimum is near to, or on, the boundary of this region. This encourages frequent activation of the reflection mechanism, allowing its performance to be effectively evaluated.

The presented experiments are implemented in Python, using the PyTorch library, and the code is available at <https://github.com/piers-hinds/RMVSDE-numeric>.

4.7.1 Approximation of the interacting particle system

To implement optimization or sampling methods based on mean-field SDEs, we need to be able to simulate the interacting particle system (4.19). Consider a uniform discretization of the time-interval $[0, T]$, for a fixed $T > 0$, such that $t_{k+1} - t_k = h$, $k = 0, \dots, K - 1$ and $T = Kh$.

In the case of CBO models like the ones from Section 4.3.1 we need numerical methods for (4.19) converging in the almost sure sense. There are two types of mean-square/almost sure methods for reflected SDEs in the literature: penalty and projection (see e.g. [109, 126]). Denote by $\Pi(x)$ the projection of x on ∂G if $x \notin \bar{G}$ else $\Pi(x) = x$. Introduce the function (penalty) $\pi(x) = (x - \Pi(x))$, which is half of the gradient of square of distance function of x from ∂G . Let $Y_k^{i,N}$ be an approximation of (4.19). The penalty and projection schemes for reflected SDEs adapted to the interacting particle system (4.19) take the form

(i) Penalty scheme :

$$\begin{aligned} Y_{k+1}^{i,N} &= Y_k^{i,N} + b(t_k, Y_k^{i,N}, \hat{\mu}_{Y_k^N})h + \sigma(t_k, Y_k^{i,N}, \hat{\mu}_{Y_k^N})\Delta W^i(t_k) \\ &\quad - \frac{h}{\epsilon}\pi(Y_k^{i,N}), \quad i = 1, \dots, N, \end{aligned} \quad (4.131)$$

where $\Delta W^i(t_k) = W^i(t_{k+1}) - W^i(t_k)$ are the Wiener increments. The optimal choice of the penalty strength ϵ is h (see [126]).

(ii) Projection scheme :

$$\begin{aligned}\bar{Y}_{k+1}^{i,N} &= Y_k^{i,N} + b(t_k, Y_k^{i,N}, \hat{\mu}_{\mathbf{Y}_k^N})h + \sigma(t_k, Y_k^{i,N}, \hat{\mu}_{\mathbf{Y}_k^N})\Delta W^i(t_k), \\ Y_{k+1}^{i,N} &= \Pi(\bar{Y}_{k+1}^{i,N}), \quad i = 1, \dots, N.\end{aligned}\tag{4.132}$$

There are rather limited numerical analysis results for mean-square convergence of these two methods. In the case of convex polyhedra, the mean-square order of convergence of both methods is $1/2$ up to a logarithmic correction, and in the case of smooth convex domains, mean-square order of $1/4$ has been proved (see, e.g., [109, 126]). Like in [71] one can potentially show that mean-square convergence of these methods is uniform in the number of particles N , considering this aspect is beyond the scope of this paper.

For sampling methods, it is sufficient to use weak-sense schemes for approximating (4.19). Weak methods for reflected SDEs are available in [96, 19, 88, 101].

4.7.2 Ackley function

We begin with the Ackley function, which is a widely used benchmark function for global optimization problems [5]. The function is defined as:

$$\begin{aligned}f(x, y) &= -20 \exp\left(-0.2\sqrt{\frac{1}{2}(x^2 + y^2)}\right) - \exp\left(\frac{1}{2}(\cos(2\pi x) + \cos(2\pi y))\right) \\ &\quad + 20 + e.\end{aligned}\tag{4.133}$$

We translate the function so that the global minimum is located at $(2, 2)$ instead of $(0, 0)$. The feasible region is the closed ball of radius 3 centred at the origin:

$$\bar{G} = \{(x, y) : x^2 + y^2 \leq 9\}.\tag{4.134}$$

The function with constraints is illustrated in Figure 3. We choose drift $\beta = 1$ and diffusion $\sigma = 4$, time $t = 1$, $\alpha = 10^4$. The initial position of the particles is sampled from a uniform distribution on the feasible

region. Both the penalty scheme (4.131) and the projection scheme (4.132) are tested and the success rates for various number of particles, N , and number of time steps, $1/h$ (in other words, choice of h with the fixed $t = 1$), are reported in Table 12. We observe that both penalty and projection schemes converge reliably to the global minimizer as N and the number of steps increase. Both approaches exhibit broadly similar performance, achieving near 100% success for modestly large N and sufficiently many steps. The projection method performs slightly better, which could be due to the fact that when the consensus is computed, all particles are guaranteed to lie within the feasible region.

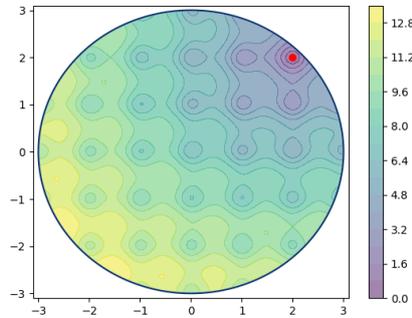


Figure 3: Translated Ackley function with minimum at $(2, 2)$ constrained to a closed ball of radius 3.

Table 12: Translated Ackley function. Comparison of success rates for the penalty and projection schemes.

(a) Penalty scheme					(b) Projection scheme				
$1/h$	N				$1/h$	N			
	10	20	50	100		10	20	50	100
5	0.055	0.137	0.416	0.717	5	0.137	0.370	0.809	0.979
10	0.256	0.605	0.966	1.000	10	0.550	0.926	1.000	1.000
20	0.762	0.986	1.000	1.000	20	0.883	0.997	1.000	1.000
50	0.725	0.973	1.000	1.000	50	0.730	0.978	1.000	1.000
100	0.317	0.728	0.993	1.000	100	0.307	0.717	0.993	1.000

4.7.3 Non-convex function with heart-shaped constraint

We now consider a non-convex function from [135] defined as

$$f(x, y) = -[\cos((x - 0.1)y)]^2 - x \sin(3x + y) \quad (4.135)$$

constrained to a heart-shaped region given by the inequality:

$$x^2 + y^2 \leq \left[2 \cos(t) - \frac{1}{2} \cos(2t) - \frac{1}{4} \cos(3t) - \frac{1}{8} \cos(4t) \right]^2 + 4 \sin^2(t), \quad (4.136)$$

where $t = \text{atan2}(x, y)$. The function and the feasible region are shown in Figure 4. The level-set nature of the constraint means that we can project particles on the boundary of the domain in the following manner. First, if a particle lies outside the feasible region, the inward normal direction to the level set can be computed. Second, using this normal direction, we can solve - by Newton's method - for the distance along this normal we need to translate the particle such that it will lie on the boundary. Note that the domain is non-convex.

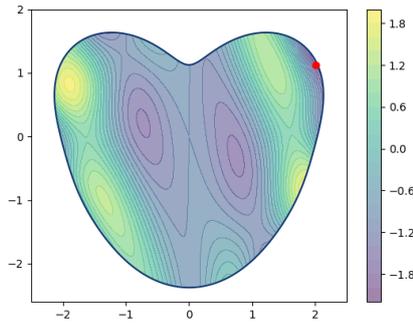


Figure 4: Non-convex function with heart-shaped constraint.

Like before, we use drift $\beta = 1$, diffusion $\sigma = 4$ and $\alpha = 10^4$. Further, we use a fixed time step of $h = 1/20$. For the initial state of the particles, we sample uniformly in a ball of radius $5/2$ before projecting onto the feasible region. We report success rates using the projection scheme (4.132) for various values of N and number of time steps K in Table 13. The results in Table 13 indicate that the CBO method reliably finds the global minimizer despite the non-convex nature of the domain. The observed success rates generally increase with N and time (number of steps K), as would be expected.

Table 13: Success rates of the CBO method on the non-convex function with heart-shaped constraint.

K	N			
	10	20	50	100
5	0.29	0.54	0.90	0.95
10	0.46	0.69	0.88	0.99
20	0.54	0.81	0.95	1.00
50	0.59	0.78	0.98	0.99
100	0.57	0.81	0.97	1.00

4.7.4 High-dimensional Rastrigin function

We consider the Rastrigin function in various dimensions. It is defined as

$$f(x) = 10d + \sum_{i=1}^d (x_i^2 - 10 \cos(2\pi x_i)), \quad (4.137)$$

constrained to the closed ball of radius 5 centred at the origin: $\bar{G} = \{x \in \mathbb{R}^d : |x| \leq 5\}$. Instead of constant drift and diffusion coefficients, β and σ , we use $\beta(t) = 10t$, $\sigma(t) = 10e^{-t \log 10}$, so that the drift increases by an order of magnitude over the interval $[0, 1]$, while σ decreases by an order of magnitude. We take $\alpha = 10^4$ and a step size of $h = 1/500$ is used. The initial position of the particles is sampled from a uniform distribution on the feasible region.

Success rates using the projection scheme (4.132) are given in Table 14 for various values of dimension, d , and number of particles, N . We also consider the cases where we have the number of iterations $K = 200$, 500, and 1000. The results in Table 14 highlight the scalability of the algorithm to higher-dimensional settings. While success rates decline with increasing dimensionality, they remain high for modestly large N . We also see an increase in success rates as K increases, particularly between $K = 200$ and $K = 500$.

Table 14: Comparison of success rates for the Rastrigin function in various dimensions for number of iterations $K = 200, 500,$ and 1000 .

(a) $K = 200$

d	N			
	10	20	50	100
5	0.194	0.472	0.834	0.982
20	0.080	0.203	0.549	0.800
100	0.031	0.086	0.237	0.417
500	0.030	0.047	0.075	0.141

(b) $K = 500$

d	N			
	10	20	50	100
5	0.251	0.533	0.899	0.993
20	0.159	0.459	0.841	0.948
100	0.095	0.275	0.750	0.950
500	0.059	0.174	0.492	0.819

(c) $K = 1000$

d	N			
	10	20	50	100
5	0.263	0.523	0.903	0.991
20	0.167	0.456	0.825	0.951
100	0.077	0.321	0.779	0.958
500	0.060	0.181	0.506	0.845

4.7.5 Testing repelling CBO on Rosenbrock function

Let us now consider the Rosenbrock function in two dimensions, defined as

$$f(x, y) = (1 - x)^2 + 100(y - x^2)^2. \quad (4.138)$$

In addition we consider the feasible region

$$\bar{G} = \{(x, y) \in \mathbb{R}^2 : x^2 + y^2 \leq 2\}, \quad (4.139)$$

such that the minimizer of f , $(1, 1)$, lies on the boundary of G . The feasible region and the Rosenbrock function are displayed in Figure 5.

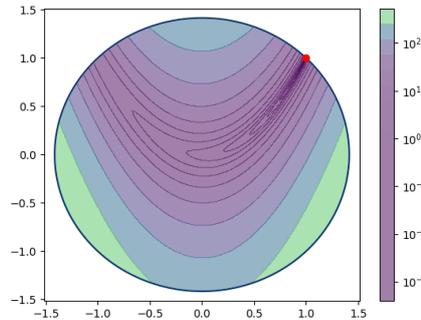


Figure 5: Rosenbrock function constrained to a closed ball of radius $\sqrt{2}$.

In this experiment, we compare the CBO method given by Equation (4.24) with the repelling CBO method (Equation (4.28)). In Table 15, success rates for both methods are shown for different values of N and K . We use the projection scheme (4.132) with a fixed step size of $h = 1/20$ for each experiment, as well as a drift parameter $\beta = 1$ and diffusion parameter $\sigma = 4$ and $\alpha = 10^4$. The initial position of the particles is sampled from a uniform distribution on the feasible region. The comparison in Table 15 illustrates the better performance of the repelling CBO method over the standard CBO method, especially for smaller N and K . The repelling method's ability to aid with exploration seems particularly relevant in this example, since it allows particles to continue exploring even when they are in a good location relative to other particles.

We remark that although the repelling method outperforms the standard CBO method, this is at the cost of additional computational complexity. Since pairwise interactions are computed, the repelling method is of $\mathcal{O}(N^2)$ complexity. However, it still requires only N objective function evaluations per iteration. Thus, we can achieve better performance without any additional function evaluations which may be relevant when, for example, the objective function is expensive to evaluate. There are ways to overcome the $\mathcal{O}(N^2)$ complexity via locality sensitive hashing or efficient nearest neighbour search, see [47] for example.

Table 15: Constrained Rosenbrock function. Comparison of success rates between standard and repelling constrained CBO.

(a) Success rates for the standard CBO

K	N			
	10	20	50	100
5	0.075	0.106	0.192	0.332
10	0.084	0.135	0.281	0.525
20	0.094	0.169	0.422	0.766
50	0.121	0.249	0.702	0.982
100	0.133	0.315	0.892	0.999

(b) Success rates for the repelling CBO

K	N			
	10	20	50	100
5	0.182	0.340	0.665	0.863
10	0.214	0.384	0.669	0.834
20	0.234	0.433	0.765	0.915
50	0.263	0.538	0.929	0.999
100	0.287	0.595	0.979	1.000

4.7.6 Inverse problem

In this section, we test the CBO model (Equation (4.24)) on an inverse problem. Consider the Cauchy problem for the parabolic partial integral differential equation (PIDE):

$$\begin{aligned} \frac{\partial u}{\partial t} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 u}{\partial x^2} - bx \frac{\partial u}{\partial x} \\ + \frac{1}{\sqrt{2\pi\gamma^2}} \int_{\mathbb{R}} [u(t, xe^y) - u(t, x)] \exp\left(-\frac{(y-m)^2}{2\gamma^2}\right) dy = 0, \end{aligned} \quad (4.140)$$

with the terminal condition

$$u(T, x) = \max\{x - 1, 0\} \quad (4.141)$$

and $b := e^{m+\frac{1}{2}\gamma^2} - 1$.

The inverse problem is formulated within Tikhonov's regularization setup as follows. Given noisy observations $\hat{u}_{i,j}$ of the solution $u(t_i, x_j)$ to the PIDE problem (4.140) at some t_i, x_j , we aim to find estimates $\hat{\theta} := (\hat{\sigma}, \hat{m}, \hat{\gamma})$ of the parameters $\theta = (\sigma, m, \gamma)$ with the constraints $\hat{\sigma} \in [0, 1]$, $\hat{m} \in [-1, 1]$, and $\hat{\gamma} \in [0, 1]$. To achieve this aim, we construct the loss function

$$\text{Loss}(\theta) = \sum_{i=1}^{10} \sum_{j=1}^5 |u(t_i, x_j; \theta) - \hat{u}_{i,j}|^2 + \lambda |\theta|_2, \quad (4.142)$$

where $\lambda > 0$ is a small regularization parameter (we use $\lambda = 10^{-5}$ in the experiments). The data is synthetically generated using the following representation for the forward map [94]:

$$\begin{aligned} u(t, x; \theta) = \sum_{j=0}^{\infty} \frac{(\tau)^j}{j!} e^{-\tau} \left\{ x e^{-b\tau + j(m+\gamma^2/2)} \Phi\left(\frac{\ln(x) + (\frac{\sigma^2}{2} - b)\tau + j(m + \gamma^2)}{\sqrt{\sigma^2\tau + j\gamma^2}}\right) \right. \\ \left. - \Phi\left(\frac{\ln(x) + (-\frac{\sigma^2}{2} - b)\tau + jm}{\sqrt{\sigma^2\tau + j\gamma^2}}\right) \right\}, \quad \tau = T - t, \end{aligned} \quad (4.143)$$

followed by adding a small amount of relative (observational) noise

$$\hat{u}_{ij} = u(t_i, x_j; \theta) + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, 10^{-3}u(t_i, x_j; \theta)). \quad (4.144)$$

Here $\Phi(\cdot)$ is the cdf for the standard normal distribution. We approximate the infinite series in (4.143) by summing the first 20 terms or truncating when the magnitude of the summand falls below 10^{-14} , whichever occurs first. Note that when the forward map is not available in an analytical form, linear PIDE problems can be solved using the Monte Carlo technique together with approximating the corresponding stochastic characteristics by a suitable numerical scheme (see e.g. [36] and references therein).

For the experiments we take $T = 3$ and use the uniformly spaced grid $t_i = 0.3(i - 1)$, $i = 1, \dots, 10$, and $x_j = 0.8 + 0.1(j - 1)$, $j = 1, \dots, 5$. As the ground truth, we chose $\sigma = 0.1$, $m = -0.2$, $\gamma = 0.3$, which are the true values to be recovered. We use $N = 400$ particles and the projection scheme (4.132) with step-size $h = 0.01$ together with the drift and diffusion coefficients described in Section 4.7.4. The initial distribution for the particles is taken to be uniform on $[0, 1] \times [-1, 1] \times [0, 1]$. We run the projection scheme for $K = 100$ steps. Additionally, we use the tolerance $\varepsilon = 0.01$ to determine if an experiment is successful. Using a value of $\alpha = 10^{14}$, we obtain a success rate of 0.99 on 100 experiments. We note that in this case, a value of $\alpha = 10^4$ is not sufficient and results in a success rate of 0.07.

In Figure 6, we include the histograms for each parameter in the final consensus of each experiment. We see that each estimated parameter is distributed reasonably tightly around the true parameter value. In a future work it is of interest to test the constrained CBO models from Section 4.3.1 on more complex inverse problems.

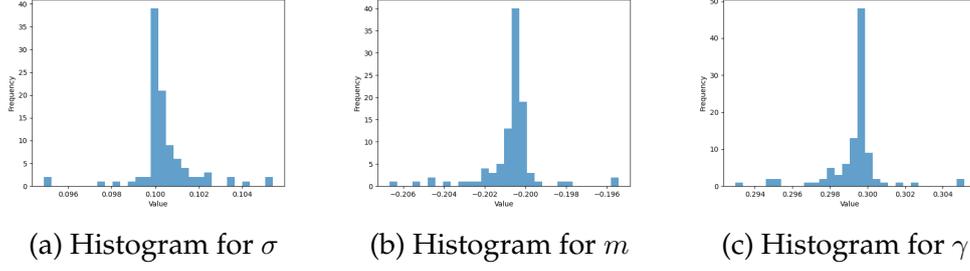


Figure 6: The inverse problem. Histograms of recovered parameters across the 100 experiments.

4.7.7 Non-Lipschitz function

We now consider an example of a non-Lipschitz function to optimize. Note that the analysis in Section 4.4 makes use of the fact that the objective function should be Lipschitz continuous (Assumption 4.2).

The (two-dimensional) Eggholder function (see [145]) is defined by

$$f(x, y) = -(y + 47) \sin\left(\sqrt{\left|\frac{x}{2} + (y + 47)\right|}\right) - x \sin\left(\sqrt{|x - (y + 47)|}\right), \quad (4.145)$$

and attains its global minimum

$$f_{\min} \approx -959.6407$$

at

$$(x^*, y^*) = (512, 404.2319).$$

We restrict the feasible space to the closed ball

$$\bar{G} = \{(x, y) : x^2 + y^2 \leq r^2\}, \quad r = \sqrt{512^2 + 404.2319^2} \approx 652.34,$$

so that $(x^*, y^*) \in \partial G$. Figure 7 shows the Eggholder function and the feasible set, as well as the true minimum.

The projection scheme is employed and an experiment is declared successful when the final consensus lies within $\varepsilon = 10$ of (x^*, y^*) . This corresponds to approximately 0.02% of the total feasible region. Drift and diffusion are chosen constant, $\beta = 10$ and $\sigma = 40$, while $\alpha = 10^3$.

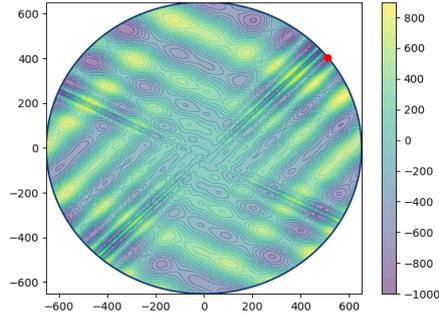


Figure 7: Eggholder function constrained to a closed ball.

Note that the increased values of ε and coefficients β and σ are due to the increased area of the feasible region.

Table 16: Eggholder function: success rates for the projection scheme.

K	N			
	10	20	50	100
5	0.136	0.208	0.411	0.548
10	0.223	0.314	0.587	0.750
20	0.319	0.498	0.768	0.889
50	0.566	0.751	0.962	0.999
100	0.753	0.945	1.000	1.000

Table 16 shows the success rates for various values of K and N . Success probability increases monotonically with the number of particles and the number of iterations. These findings confirm that the projection-based CBO algorithm is practical even for the highly oscillatory, non-Lipschitz, Eggholder function.

4.8 Conclusion and future research

This chapter was concerned with reflected McKean-Vlasov SDEs in smooth bounded domains and their application to modelling consensus-based optimization algorithms. We established well-posedness of general reflected McKean-Vlasov SDEs under some mild assumptions on the coefficients. We then proved the propagation of chaos result with the optimal rate of convergence. Two consensus-based optimization models were proposed, and their coefficients are shown to satisfy the conditions for

well-posedness and propagation of chaos. Under a Lipschitz condition on the objective function, we showed convergence of the mean-field limit to the optimizer of the objective. Finally, we considered two numerical schemes for the CBO models, which correspond to optimization algorithms. These algorithms were tested on several standard benchmark problems from constrained optimization and shown to be effective.

Let us now outline some directions for future research. A natural extension of this work is to study exit times in the finite- N particle approximation of the constrained CBO scheme, and in its mean-field limit. In stochastic optimization, analogous questions have been studied for stochastic gradient descent and Langevin-type algorithms, where one analyses the time required for the continuous dynamics to escape a neighbourhood of a local minimum due to stochastic perturbations (see, e.g., [137, 106]). For the interacting CBO system, a similar notion of escape can be formulated in terms of the consensus or empirical measure leaving a basin of attraction. Recent progress on reflected McKean–Vlasov equations [1] establishes a large-deviation framework for such systems, providing asymptotic estimates for exit probabilities and times under reflection constraints. Extending these results to the constrained CBO setting would enable an understanding of how noise intensity and the number of particles jointly determine the ability of the system to leave suboptimal regions of the objective landscape.

There has also been work on the inclusion of momentum into the dynamics of CBO [32]. Momentum plays a positive role in many gradient-based algorithms, both theoretically and practically [113, 103, 78] and indeed it is shown to be effective in the context of CBO in [32]. Incorporating momentum into constrained CBO algorithms is a natural extension, although theoretically it is not completely straightforward due to the local time term in the SDEs (see e.g. [89]). Some preliminary numerical experiments have shown this to be a promising direction of future research.

5 Conclusion

In this thesis, we have explored two distinct topics within stochastic numerics: variance reduction methods for Monte Carlo integration of stochastic differential equations, and the study of reflected McKean–Vlasov SDEs in bounded domains.

In the first area, we developed a variance reduction technique utilizing neural networks to approximate optimal control variates, particularly for Lévy-driven SDEs. We established the theoretical conditions under which perfect variance reduction could be achieved, and developed an algorithm which significantly enhances computational efficiency. Numerical examples from option pricing illustrated the practical efficacy of our proposed methods, providing substantial improvements over standard Monte Carlo approaches.

The second focus of our work addressed the well-posedness and numerical approximation of reflected McKean–Vlasov SDEs, with applications in constrained optimization and consensus-based methods. We established theoretical results confirming well-posedness, and we provided optimal convergence rates for particle approximations. Two consensus-based models were considered and shown to be practically effective.

Together, these studies extend the theoretical understanding of the respective topic and provide practical methods for solving problems in finance and optimization. Future research directions were considered in the concluding sections 3.6 and 4.8.

References

- [1] Adams, D., Reis, G. d., Ravaille, R., Salkeld, W., and Tugaut, J. “Large Deviations and Exit-times for reflected McKean–Vlasov equations with self-stabilising terms and superlinear drifts”. *Stochastic Processes and their Applications* 146 (2022), pp. 264–310.
- [2] Allen, L. J. *An introduction to stochastic processes with applications to biology*. CRC press, 2010.
- [3] Applebaum, D. *Lévy Processes and Stochastic Calculus*. 2nd ed. Cambridge: Cambridge University Press, 2009.
- [4] Asmussen, S. and Rosiński, J. “Approximations of small jumps of Lévy processes with a view towards simulation”. *J. Appl. Probab.* 38(2) (2001), pp. 482–493.
- [5] Back, T. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford Univ. Press, 1996.
- [6] Bae, H.-O., Ha, S.-Y., Kang, M., Lim, H., Min, C., and Yoo, J. “A constrained consensus based optimization algorithm and its application to finance”. *Applied Mathematics and Computation* 416 (2022).
- [7] Barron, A. R. “Approximation and estimation bounds for artificial neural networks”. *Machine Learning* 14(1) (1994), pp. 115–133.
- [8] Beddrich, J., Chenchene, E., Fornasier, M., Huang, H., and Wohlmuth, B. *Constrained consensus-based optimization and numerical heuristics for the few particle regime*. arXiv: 2410.10361, 2024.
- [9] Bellomo, N., Burini, D., Dosi, G., Gibelli, L., Knopoff, D., Outada, N., Terna, P., and Virgillito, M. E. “What is life? A perspective of the mathematical kinetic theory of active particles”. *Math. Models Methods Appl. Sci.* 31(09) (2021), pp. 1821–1866.

- [10] Bellomo, N., Liao, J., Quaini, A., Russo, L., and Siettos, C. “Human behavioral crowds review, critical analysis and research perspectives”. *Math. Models Methods Appl. Sci.* 33(08) (2023), pp. 1611–1659.
- [11] Belomestny, D., Milstein, G. N., and Schoenmakers, J. G. M. “Sensitivities for Bermudan options by regression methods”. *Decisions Econ. Finance* 33 (2010), pp. 117–138.
- [12] Belomestny, D., Häfner, S., Nagapetyan, T., and Urusov, M. “Variance reduction for discretised diffusions via regression”. *J. Mathem. Anal. Applic.* 458(1) (2018), pp. 393–418.
- [13] Berner, J., Grohs, P., and Jentzen, A. “Analysis of the generalization error: Empirical risk minimization over deep artificial neural networks overcomes the curse of dimensionality in the numerical approximation of Black–Scholes partial differential equations”. *SIAM J. Mathem. Data Science* 2(3) (2020), pp. 631–657.
- [14] Berner, J., Grohs, P., Kutyniok, G., and Petersen, P. “The modern mathematics of deep learning”. *Mathematical Aspects of Deep Learning*. Ed. by Grohs, P. and Kutyniok, G. Cambridge University Press, 2022, pp. 1–111.
- [15] Billingsley, P. *Probability and measure*. John Wiley & Sons, 2017.
- [16] Bogachev, V. I. and Kolesnikov, A. V. “The Monge-Kantorovich problem: achievements, connections, and perspectives”. en. *Russian Math. Surveys* 67(5) (2012), pp. 785–890.
- [17] Bolley, F. “Separability and completeness for the Wasserstein distance”. en. *Séminaire de Probabilités XLI*. Ed. by Donati-Martin, C., Émery, M., Rouault, A., and Stricker, C. Lecture Notes in Mathematics. Springer, 2008, pp. 371–377.
- [18] Borghi, G., Herty, M., and Pareschi, L. “Constrained consensus-based optimization”. *SIAM Journal on Optimization* 33(1) (2023), pp. 211–236.

- [19] Bossy, M., Gobet, E., and Talay, D. “A symmetrized Euler scheme for an efficient approximation of reflected diffusions”. *J. Appl. Probab.* 41(3) (2004), pp. 877–889.
- [20] Bossy, M. and Talay, D. “A stochastic particle method for the McKean-Vlasov and the Burgers equation”. *Mathematics of computation* 66(217) (1997), pp. 157–192.
- [21] Bouchard, B. “Optimal reflection of diffusions and barrier options pricing under constraints”. *SIAM J. Control Optim.* 47(4) (2008), pp. 1785–1813.
- [22] Boyd, S. P. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- [23] Brigo, D. and Mercurio, F. *Interest rate models – theory and practice: with smile, inflation and credit*. Berlin: Springer, 2006.
- [24] Caflisch, R. E. “Monte carlo and quasi-monte carlo methods”. *Acta numerica* 7 (1998), pp. 1–49.
- [25] Carmona, R. and Delarue, F. *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.
- [26] Carmona, R., Delarue, F., and Lacker, D. “Mean field games with common noise”. *Annals of Probability* 44(6) (2016), pp. 3740–3803.
- [27] Carmona, R., Fouque, J.-P., and Sun, L.-H. “Mean Field Games and systemic risk”. *Communications in Mathematical Sciences* 13(4) (2015), pp. 911–933.
- [28] Carrillo, J. A., Choi, Y.-P., Totzeck, C., and Tse, O. “An analytical framework for consensus-based global optimization method”. *Math. Models Methods Appl. Sci.* 28(06) (2018), pp. 1037–1066.
- [29] Carrillo, J. A., Clini, A., and Solem, S. “The mean field limit of stochastic differential equation systems modeling grid cells”. *SIAM J. Math. Anal.* 55(4) (2023), pp. 3602–3634.

- [30] Carrillo, J. A., Jin, S., Li, L., and Zhu, Y. "A consensus-based global optimization method for high dimensional machine learning problems". *ESAIM. COCV* 27 (2021), S5.
- [31] Carrillo, J. A., Totzeck, C., and Vaes, U. "Consensus-based optimization and ensemble Kalman inversion for global optimization problems with constraints". *Modeling and Simulation for Collective Dynamics*. World Scientific, 2023, pp. 195–230.
- [32] Chen, J. and Lyu, L. "A Consensus-Based Global Optimization Method with Adaptive Momentum Estimation". *Communications in Computational Physics* 31(4) (2022), pp. 1296–1316.
- [33] Chen, X., Dos Reis, G., and Stockinger, W. "Wellposedness, exponential ergodicity and numerical approximation of fully super-linear McKean–Vlasov SDEs and associated particle systems". *Electronic Journal of Probability* 30 (2025), pp. 1–50.
- [34] Cont, R. and Tankov, P. *Financial modelling with jump processes*. Boca Raton: Chapman & Hall/CRC, 2004.
- [35] Cybenko, G. "Approximation by superpositions of a sigmoidal function". *Mathematics of Control, Signals, and Systems* 2(4) (1989), pp. 303–314.
- [36] Deligiannidis, G., Maurer, S., and Tretyakov, M. V. "Random walk algorithm for the Dirichlet problem for parabolic integro-differential equation". *BIT Numer. Math.* 61 (2021), pp. 1223–1269.
- [37] Dembo, A. and Zeitouni, O. *Large deviations techniques and applications*. Vol. 38. Springer Science & Business Media, 2009.
- [38] Dick, J., Kuo, F., and Sloan, I. "High-dimensional integration: the quasi-Monte Carlo way". *Acta Numerica* 22 (2013), pp. 133–288.
- [39] Dinguç, K. D. and Hörmann, W. "A general control variate method for option pricing under Lévy processes". *European J. Oper. Research* 221(2) (2012), pp. 368–377.

- [40] Dupuis, P. and Ishii, H. “On Lipschitz continuity of the solution mapping to the Skorokhod problem, with applications”. *Stochastics: An International Journal of Probability and Stochastic Processes* 35(1) (1991), pp. 31–62.
- [41] Dupuis, P. and Ishii, H. “SDEs with oblique reflection on nonsmooth domains”. *The annals of Probability* (1993), pp. 554–580.
- [42] Dynkin, E. B. *Markov processes*. Berlin: Springer, 1965.
- [43] Fornasier, M., Klock, T., and Riedl, K. “Consensus-based optimization methods converge globally”. *arXiv preprint arXiv:2103.15130* (2021).
- [44] Fournié, E., Lasry, J.-M., Lebuchoux, J., Lions, P.-L., and Touzi, N. “Applications of Malliavin calculus to Monte Carlo methods in finance”. *Finance and Stochastics* 3(4) (1999), pp. 391–412.
- [45] Fournier, N. and Guillin, A. “On the rate of convergence in Wasserstein distance of the empirical measure”. *Probab. Theory Related Fields* 162(3) (2015), pp. 707–738.
- [46] Freidlin, M. I. *Functional integration and partial differential equations*. Princeton: Princeton Univ. Press, 1985.
- [47] Frenkel, D. and Smit, B. *Understanding molecular simulation: from algorithms to applications*. Elsevier, 2023.
- [48] Friedman, A. *Stochastic differential equations and applications*. Berlin: Springer, 2011.
- [49] Garbuno-Inigo, A., Hoffmann, F., Li, W., and Stuart, A. M. “Interacting Langevin diffusions: gradient structure and ensemble Kalman sampler”. *SIAM J. Appl. Dynam. Syst.* 19(1) (2020), pp. 412–441.
- [50] Garbuno-Inigo, A., Nusken, N., and Reich, S. “Affine invariant interacting Langevin dynamics for Bayesian inference”. *SIAM J. Appl. Dynam. Syst.* 19(3) (2020), pp. 1633–1658.

- [51] Garroni, M. G. and Menaldi, J. L. *Second order elliptic integro-differential problems*. New York: Chapman & Hall/CRC, 2002.
- [52] Geist, M., Petersen, P., Raslan, M., Schneider, R., and Kutyniok, G. “Numerical solution of the parametric diffusion equation by deep neural networks”. *J. Scien. Comp.* 88 (2021).
- [53] Gikhman, I. I. and Skorokhod, A. V. *The theory of stochastic processes II*. Springer Science & Business Media, 2004.
- [54] Gikhman, I. I. and Skorokhod, A. V. *Stochastic differential equations*. Springer, 2007.
- [55] Giles, M. B. “Multilevel Monte Carlo methods”. *Acta Numerica* 24 (2015), pp. 259–328.
- [56] Gladyshev, S. A. and Milstein, G. N. “The Runge–Kutta method for calculation of Wiener integrals of functionals of exponential type”. *Zh. Vychisl. Mat. i Mat. Fiz.* 24 (1984), pp. 1136–1149.
- [57] Glasserman, P. *Monte Carlo methods in financial engineering*. Berlin: Springer, 2003.
- [58] Gonon, L., Grohs, P., Jentzen, A., Kofler, D., and Šiška, D. “Uniform error estimates for artificial neural network approximations for heat equations”. *IMA J. Numer. Anal.* 42(3) (2022), pp. 1991–2054.
- [59] Goodfellow, I., Bengio, Y., and Courville, A. *Deep learning*. MIT press, 2016.
- [60] Granas, A., Dugundji, J., et al. *Fixed point theory*. Vol. 14. Springer, 2003.
- [61] Ha, S.-Y., Jin, S., and Kim, D. “Convergence of a first-order consensus-based global optimization algorithm”. *Math. Models Methods Appl. Sci.* 30(12) (2020), pp. 2417–2444.
- [62] Ha, S.-Y. and Tadmor, E. “From particle to kinetic and hydrodynamic descriptions of flocking”. *Kinetic and Related Models* 1(3) (2008), pp. 415–435.

- [63] Han, J., Jentzen, A., and E, W. “Solving high-dimensional partial differential equations using deep learning”. *Proc. National Acad. Sci.* 115(34) (2018), pp. 8505–8510.
- [64] Heston, S. L. “A closed-form solution for options with stochastic volatility with applications to bond and currency options”. *Review Finan. Studies* 6 (1993), pp. 327–343.
- [65] Hinds, P. D., Sharma, A., and Tretyakov, M. V. “Well-posedness and approximation of reflected McKean-Vlasov SDEs with applications”. *Mathematical Models and Methods in Applied Sciences* 35(08) (2025), pp. 1845–1887.
- [66] Hinds, P. D. and Tretyakov, M. V. “Neural variance reduction for stochastic differential equations”. *Journal of Computational Finance* 27(3) (2023), pp. 1–41.
- [67] Hodgkinson, L., Heide, C. van der, Roosta, F., and Mahoney, M. W. “Stochastic continuous normalizing flows: training SDEs as ODEs”. *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*. Ed. by Campos, C. de and Maathuis, M. H. Vol. 161. Proceedings of Machine Learning Research. PMLR, 2021, pp. 1130–1140.
- [68] Hornik, K., Stinchcombe, M., and White, H. “Multilayer feedforward networks are universal approximators”. *Neural Networks* 2(5) (1989), pp. 359–366.
- [69] Horst, R., Pardalos, P. M., and Van Thoai, N. *Introduction to global optimization*. Springer Science & Business Media, 2000.
- [70] Kac, M. “Foundations of kinetic theory”. *Proceedings of The third Berkeley symposium on mathematical statistics and probability*. Vol. 3. 600. 1956, pp. 171–197.
- [71] Kalise, D., Sharma, A., and Tretyakov, M. V. “Consensus-based optimization via jump-diffusion stochastic differential equations”. *Mathematical Models and Methods in Applied Sciences* 33(2) (2023), pp. 289–339.

- [72] Kampen, N. G. van. *Stochastic processes in physics and chemistry*. 3rd ed. Amsterdam: North Holland, 2007.
- [73] Karatzas, I. and Shreve, S. *Brownian motion and stochastic calculus*. Vol. 113. Springer Science & Business Media, 1991.
- [74] Ken-Iti, S. *Lévy processes and infinitely divisible distributions*. Vol. 68. Cambridge university press, 1999.
- [75] Kennedy, J. and Eberhart, R. "Particle swarm optimization". Vol. 4. IEEE, 1995, pp. 1942–1948.
- [76] Khasminskii, R. *Stochastic stability of differential equations*. 2nd ed. Springer Berlin, Heidelberg, 2012.
- [77] Kidger, P., Foster, J., Li, X., and Lyons, T. J. "Neural SDEs as infinite-dimensional GANs". *Intern. Confer. Machine Learning*. Vol. 139. PMLR, 2021, pp. 5453–5463.
- [78] Kingma, D. P. and Ba, J. "Adam: a method for stochastic optimization". *Intern. Confer. Learning Representations (2015)*.
- [79] Klebaner, F. C. *Introduction to stochastic calculus with applications*. World Scientific Publishing Company, 2012.
- [80] Kloeden, P. E. and Platen, E. *Numerical solution of stochastic differential equations*. Springer, 1992.
- [81] Kohatsu-Higa, A., Ortiz-Latorre, S., and Tankov, P. "Optimal simulation schemes for Lévy driven stochastic differential equations". *Math. Comp.* 83(289) (2014), pp. 2293–2324.
- [82] Kou, S. G. "A jump-diffusion model for option pricing". *Management science* 48(8) (2002), pp. 1086–1101.
- [83] Kreyszig, E. *Introductory functional analysis with applications*. John Wiley & Sons, 1991.
- [84] Kyprianou, A. E. *Introductory lectures on fluctuations of Lévy processes with applications*. Springer, 2006.

- [85] Lacker, D. “On a Strong Form of Propagation of Chaos for McKean-Vlasov Equations”. *Electronic Communications in Probability* 23 (2018), pp. 1–11.
- [86] Lasry, J.-M. and Lions, P.-L. “Mean field games”. *Japanese journal of mathematics* 2(1) (2007), pp. 229–260.
- [87] Leimkuhler, B., Matthews, C., and Weare, J. “Ensemble preconditioning for Markov chain Monte Carlo simulation”. *Stat. Comput.* 28 (2018), pp. 277–290.
- [88] Leimkuhler, B., Sharma, A., and Tretyakov, M. V. “Simplest random walk for approximating Robin boundary value problems and ergodic limits of reflected diffusions”. *Ann. Appl. Probab.* 33(3) (2023), pp. 1904–1960.
- [89] Leimkuhler, B., Sharma, A., and Tretyakov, M. V. “Numerical integrators for confined Langevin dynamics”. *arXiv preprint arXiv:2404.16584* (2024).
- [90] Lions, P.-L. and Sznitman, A.-S. “Stochastic differential equations with reflecting boundary conditions”. *Communications on pure and applied Mathematics* 37(4) (1984), pp. 511–537.
- [91] Liu, Q. and Wang, D. “Stein variational gradient descent: A general purpose Bayesian inference algorithm”. *Proceedings of the 30th Intern. Conference on Neural Information Processing Systems*. NIPS’16. Barcelona, Spain: Curran Associates Inc., 2016, pp. 2378–2386.
- [92] Maruyama, G. “Continuous Markov processes and stochastic equations”. *Rendiconti del Circolo Matematico di Palermo* 4 (1955), pp. 48–90.
- [93] McKean Jr, H. P. “A class of Markov processes associated with nonlinear parabolic equations”. *Proceedings of the National Academy of Sciences* 56(6) (1966), pp. 1907–1911.
- [94] Merton, R. C. “Option pricing when underlying stock returns are discontinuous”. *J. Finan. Econom.* 3(1) (1976), pp. 125–144.

- [95] Milstein, G. N. "A method of second-order accuracy integration of stochastic differential equations". *Theory of Probability & Its Applications* 23(2) (1979), pp. 396–401.
- [96] Milstein, G. N. "Application of the numerical integration of stochastic equations for the solution of boundary value problems with Neumann boundary conditions". *Theor. Prob. Appl.* 41 (1997), pp. 170–177.
- [97] Milstein, G. N. "Weak approximation of solutions of systems of stochastic differential equations". *Theor. Probab. Appl.* 30 (1985), pp. 706–721.
- [98] Milstein, G. N. and Schoenmakers, J. G. M. "Monte Carlo construction of hedging strategies against multi-asset European claims". *Stoch. Stoch. Reports* 73(1-2) (2002), pp. 125–157.
- [99] Milstein, G. N. and Tretyakov, M. V. "Numerical analysis of Monte Carlo evaluation of Greeks by finite differences". *J. Comp. Finance* 8 (2005), pp. 1–33.
- [100] Milstein, G. N. and Tretyakov, M. V. "Practical variance reduction via regression for simulating diffusions". *SIAM J. Numer. Anal.* 47(2) (2009), pp. 887–910.
- [101] Milstein, G. N. and Tretyakov, M. V. *Stochastic numerics for mathematical physics*. 2nd Edition. Berlin: Springer, 2021.
- [102] Mishura, Y. and Veretennikov, A. "Existence and uniqueness theorems for solutions of McKean–Vlasov stochastic equations". *Theory of Probability and Mathematical Statistics* 103 (2020), pp. 59–101.
- [103] Nesterov, Y. "A method for solving the convex programming problem with convergence rate $O(1/k^2)$ ". *Dokl akad nauk Sssr*. Vol. 269. 1983, p. 543.
- [104] Newton, N. J. "Continuous-time Monte Carlo methods and variance reduction". *Numerical methods in finance*. Ed. by Rogers, L. and Talay, D. Cambridge: Cambridge University Press, 1997, pp. 22–42.

- [105] Newton, N. J. “Variance reduction for simulated diffusions”. *SIAM J. Appl. Math.* 54(6) (1994), pp. 1780–1805.
- [106] Nguyen, T. H., Simsekli, U., Gurbuzbalaban, M., and Richard, G. “First exit time analysis of stochastic gradient descent under heavy-tailed gradient noise”. *Advances in neural information processing systems* 32 (2019).
- [107] Nocedal, J. and Wright, S. *Numerical optimization*. Springer, 2006.
- [108] Øksendal, B. *Stochastic Differential Equations: An Introduction with Applications*. Universitext. Springer Berlin, Heidelberg, 2010.
- [109] Pettersson, R. “Approximations for stochastic differential equations with reflecting convex boundaries”. *Stoch. Proc. Applic.* 59 (1995), pp. 295–308.
- [110] Pilipenko, A. *An introduction to stochastic differential equations with reflection*. Lect. Pure Appl. Math. Universitätsverlag Potsdam, 2014.
- [111] Pinnau, R., Totzeck, C., Tse, O., and Martin, S. “A consensus-based model for global optimization and its mean-field limit”. *Math. Models Methods Appl. Sci.* 27(01) (2017), pp. 183–204.
- [112] Platen, E. and Bruti-Liberati, N. *Numerical solution of stochastic differential equations with jumps in finance*. Berlin: Springer, 2010.
- [113] Polyak, B. T. “Some methods of speeding up the convergence of iteration methods”. *Ussr computational mathematics and mathematical physics* 4(5) (1964), pp. 1–17.
- [114] Protter, P. E. and Protter, P. E. *Stochastic differential equations*. Springer, 2005.
- [115] Reich, S. and Cotter, C. *Probabilistic forecasting and Bayesian data assimilation*. Cambr. Univ. Press, 2015.
- [116] Reif, F. *Fundamentals of statistical and thermal physics*. Waveland Press, 2009.

- [117] Reis, G. dos, Engelhardt, S., and Smith, G. "Simulation of McKean–Vlasov SDEs with super-linear growth". *IMA Journal of Numerical Analysis* 42(1) (2022), pp. 874–922.
- [118] Rogers, L. C. G. and Williams, D. *Diffusions, Markov Processes, and Martingales*. 2nd ed. Cambridge Mathematical Library. Cambridge University Press, 2000.
- [119] Saisho, Y. "Stochastic differential equations for multi-dimensional domain with reflecting boundary". *Probability Theory and Related Fields* 74(3) (1987), pp. 455–477.
- [120] Serrin, J. "The Problem of Dirichlet for Quasilinear Elliptic Differential Equations with Many Independent Variables". *Phil. Trans. Royal Soc. London Series A* 264(1153) (1969), pp. 413–496.
- [121] Shiraya, K. and Takahashi, A. "A general control variate method for multi-dimensional SDEs: An application to multi-asset options under local stochastic volatility with jumps models in finance". *European J. Oper. Res.* 258(1) (2017), pp. 358–371.
- [122] Shiraya, K., Uenishi, H., and Yamazaki, A. "A general control variate method for Lévy models in finance". *European J. Oper. Res.* 284(3) (2020), pp. 1190–1200.
- [123] Shiryaev, A. N. *Probability-1*. Vol. 95. Springer, 2016.
- [124] Sirignano, J. and Spiliopoulos, K. "DGM: A deep learning algorithm for solving partial differential equations". *J. Comput. Phys.* 375 (2018), pp. 1339–1364.
- [125] Skorokhod, A. V. "Stochastic Equations for Diffusion Processes in a Bounded Region". *Theory of Probability & Its Applications* 6(3) (1961), pp. 264–274.
- [126] Slominski, L. "Euler's approximations of solutions of SDEs with reflecting boundary". *Stoch. Proc. Applic.* 94 (2001), pp. 317–337.

- [127] Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. “Score-Based Generative Modeling through Stochastic Differential Equations”. *International Conference on Learning Representations*.
- [128] Storn, R. and Price, K. “Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces”. *Journal of global optimization* 11 (1997), pp. 341–359.
- [129] Su, H., Tretyakov, M., and Newton, D. P. “Deep learning of transition probability densities for stochastic asset models with applications in option pricing”. *Management Science* 71(4) (2025), pp. 2922–2952.
- [130] Sznitman, A.-S. “Topics in propagation of chaos”. *Ecole d’été de probabilités de Saint-Flour XIX—1989* 1464 (1991), pp. 165–251.
- [131] Sznitman, A.-S. “Nonlinear reflecting diffusion process, and the propagation of chaos and fluctuations associated”. *Journal of functional analysis* 56(3) (1984), pp. 311–336.
- [132] Tanaka, H. “Stochastic differential equations with reflecting boundary condition in convex regions”. *Hiroshima Mathematical Journal* 9(1) (1979), pp. 163–177.
- [133] Tankov, P. and Cont, R. *Financial Modelling with Jump Processes, Second Edition*. Taylor & Francis, 2015.
- [134] Totzeck, C. “Trends in consensus-based optimization”. *Active Particles, Volume 3: Advances in Theory, Models, and Applications*. Springer, 2021, pp. 201–226.
- [135] Townsend, A. “Constrained optimization in Chebfun”. www.chebfun.org/examples/opt/Con (2014).
- [136] Tretyakov, M. V. *Introductory course on financial mathematics*. London: ICP, 2013.

- [137] Tzen, B., Liang, T., and Raginsky, M. “Local optimality and generalization guarantees for the Langevin algorithm via empirical metastability”. *Conference On Learning Theory*. PMLR. 2018, pp. 857–875.
- [138] Tzen, B. and Raginsky, M. “Neural stochastic differential equations: deep latent Gaussian models in the diffusion limit”. *arXiv:1905.09883* (2019).
- [139] Vidales, M. S., Siska, D., and Szpruch, L. “Unbiased deep solvers for linear parametric PDEs”. *Applied Mathematical Finance* 28(4) (2021), pp. 299–329.
- [140] Villani, C. et al. *Optimal transport: old and new*. Vol. 338. Springer, 2008.
- [141] Voß, S. “Meta-heuristics: The state of the art”. Springer, 2000, pp. 1–23.
- [142] Wagner, W. “Monte Carlo evaluation of functionals of solutions of stochastic differential equations. Variance reduction and numerical examples”. *Stoch. Anal. Appl.* 6 (1988), pp. 447–468.
- [143] Walter, W. *Ordinary differential equations*. Vol. 182. Springer Science & Business Media, 2013.
- [144] Wang, F.-Y. “Distribution dependent reflecting stochastic differential equations”. *Sci. China Math.* 66(11) (2023), pp. 2411–2456.
- [145] Whitley, D., Rana, S., Dzuber, J., and Mathias, K. E. “Evaluating evolutionary algorithms”. *Artificial Intelligence* 85(1) (1996), pp. 245–276.